


CONNEXIONS[®]



The Interoperability Report

July 1994 *Special Issue: Focus on Asia* Volume 8, No. 7

*ConneXions —
The Interoperability Report
tracks current and emerging
standards and technologies
within the computer and
communications industry.*

In this issue:

Internetworking in Asia.....	2
The CAE-NET project.....	7
The UNCL project.....	12
Asia Pacific NIC.....	14
IP over ATM.....	20
Book Reviews.....	26
Announcements.....	30

ConneXions is published monthly by Interop Company, a division of ZD Expos, 303 Vintage Park Drive, Foster City, California, 94404-1138, USA.
Phone: +1 (415) 578-6900
Fax: +1 (415) 525-0194
E-mail: connexions@interop.com

Subscription hotline: 1-800-575-5717
or +1-502-493-3217

Copyright © 1994 by Interop Company.
Quotation with attribution encouraged.
ConneXions—The Interoperability Report
and the *ConneXions* logo are registered
trademarks of Interop Company.

ISSN 0894-5926

From the Editor

Welcome to this special edition of *ConneXions*. This issue is being released for NetWorld+Interop 94 in Tokyo and contains articles which focus on internetworking in Asia. Since the Internet is growing rapidly in all parts of the world, we're likely to see more articles from the Asia Pacific region in future issues.

Our first article by Kilman Chon of the Korea Advanced Institute of Science and Technology is an overview of networking in Asia and a look at the work of the *Asia-Pacific Coordinating Committee for Intercontinental Research Networking* (APCCIRN). Dr. Chon is a member of the program committee for NetWorld+Interop 94 Tokyo.

Next we look at two major networking projects in Japan. Nippon Telephone and Telegraph (NTT) has been working on a distributed software development environment based on UNIX workstations and computer networks. Called CAE or *Computer Aided Engineering*, the project revolves around a nation-wide network connecting not only NTT and its affiliates, but also other companies engaged in joint development. This network, CAE-NET, provides an environment for distributed software development with reusable modules. The article is by Ryoichi Hosoya and Shunichi Fukuyama of NTT.

Gigabit networking is a major area for research in many parts of the world. In Japan, the *Ultra-high Speed Network and Computer Technology Laboratories* (UNCL) have been set up to develop an integrated system of high-speed computers and networks. The project is described here by Masahiro Taka, the director of UNCL.

The Internet has no "central administration," but relies instead on a number of cooperative distributed mechanisms for operation. In the early days of the ARPANET, a central *Network Information Center* (NIC) existed to assist both end users and network administrators. These days, NICs are operated by both local and regional network providers. National and regional NICs are generally not end-user organizations, but focus instead on address allocation and coordination. David Conrad describes the work of the *Asia Pacific Network Information Center* (APNIC).

Asynchronous Transfer Mode (ATM) is a rapidly emerging technology. We asked Mark Laubach to give us a status report on ATM as it relates to TCP/IP, and to describe the BAGNET ATM pilot project.

We want to let you know about an important change for subscribers. Customer service and fulfillment is now being handled by the Cobb Group in Louisville, Kentucky. This change means that from now on you will be receiving real subscription invoices and regular renewal letters (!) You can reach customer service directly at 1-800-575-5717 or 1-502-493-3217. Our editorial address remains the same.

Internetworking in Asia

by

Kilnam Chon,

Korea Advanced Institute of Science and Technology

Brief history

Computer networking in Asia became popular with readily available network software packages such as TCP/IP, UUCP and USENET in the early 1980s. [3] Other networks such as BITNET and FidoNet were introduced in the 1980s. Nationwide network implementation based on these network software packages were carried out among various countries in Asia throughout the 1980s. The implementation naturally led to the international links, mostly to the USA and some within Asia. By the early 1990s, most countries in Asia have domestic networks with international links. See Table 1 and Figure 1 for further information.

Coordination

The *Coordinating Committee for Intercontinental Research Networking* (CCIRN) and its regional committees in North America and Europe were formed in 1987 to coordinate networking for global research and education communities. [4] The *Asia-Pacific Coordinating Committee for Intercontinental Research Networking* (APCCIRN) was formed in the early 1990s and joined CCIRN. APCCIRN and its engineering planning group, the *Asia-Pacific Engineering Planning Group* (APEPG) meets twice a year; once during the INET conference, and once in the Asia-Pacific region.

APCCIRN started the pilot project *Asia-Pacific Network Information Center* (APNIC), which coordinates global network information center activities with InterNIC in the US and RIPE NCC in Europe. See the article on APNIC in this issue for further information on APNIC. [10]

APCCIRN set up several working groups to work on the following issues;

- Commercial Services
- Developing Countries
- Local Language Support
- Link Coordination.

Commercial service is increasingly important in the Internet, and it is expected to become the dominant service provision in the US soon and other countries later. In Asia, as a latecomer to the Internet, the commercial services started almost at the same time as the non-profit operations, unlike in North America and Europe.

There are two types of services in addition to the traditional non-profit operations for research and education communities. They are *commercial* services, and *general* services. The former are for-profit operations. The latter are not necessarily for-profit operations, but rather mechanisms for cost sharing. Many countries started the commercial and/or general services. Table 2 shows the current status compiled at APCCIRN Secretariat.

The *ap-commercial* Working Group was formed at the APCCIRN/APEPG Meeting last December, and it was headed by B. Coggershall of Supernet in Hong Kong.

Developing countries

Asia is a unique continent where the GNP per capita ranges from \$150 to \$30,000. Developed and developing countries coexist. The developed countries have fairly good domestic networking infrastructure with T1 or Fractional T1 international links.

In some developing countries, networking activities are negligible or e-mail is not available yet. Networking is available in over 130 countries, i.e., over eighty percent of the world. [6] Many of the non-networked countries are in the Asia-Pacific region and the Arab-Africa region. We need to work on these countries in addition to improve the networking capability of weakly-networked countries.

We need to train people, and support the setting up of networks and development of applications. APCCIRN set up the working group, *ap-develop* last year, which is headed by Dr. D. Narayan of Science University of Tokyo in Japan. UNESCO has been working in this area through its *Information Infrastructure Program*. [7] Regional Informatics Network Projects have been implemented as a part of the program. RINSEAP and RINSCA were established for Southeast Asia and Pacific, and South and Central Asia, respectively. Other organizations of the United Nations such as WHO, and UNDP as well as the World Bank are working to set up a global network infrastructure.

We need to come up with regional solutions to work on the non-networked and the weakly networked countries systematically. These countries may be better supported regionally by the center of excellence proposed by UNESCO. Some of regional institutions could play an instrumental role. Or, we may set up a new institution to work on this issue and others.

<i>Country/Region with International Dialup</i>	<i>Country/Region with International Leased Line</i>
Australia	French Polynesia
China	Indonesia
Fiji	Iran
Hong Kong	Pakistan
India	Sri Lanka
Japan	Vietnam
Korea	
Malaysia	
New Zealand	
Philippines	
Singapore	
Taiwan	
Thailand	

Table 1: International Connectivity

<i>Country</i>	<i>Commercial Svc Provider</i>	<i>General Svc Provider</i>
Hong Kong	Supernet	
India		Softnet
Japan	ATT-JENS, IJ, PSI Japan	
Korea	(Korea Telecom), (Dacom)	
Malaysia		(MIMOS)
Singapore		Technet
Taiwan	HiNet	Seednet

Table 2: Service Providers

Local language support

Many Asian countries use languages quite different from English with large character sets. This and other differences force us to modify the existing network software substantially. Sometimes, we are forced to develop software from scratch due to substantial difference from the English version. Local language support is a rather unique issue to Asia, and we need to pay particular attention to it. Otherwise, networking is good only in English, which limits usage, and we may end up with user-unfriendly systems.

continued on next page

Internetworking in Asia (*continued*)

APCCIRN has the working group, *ap-i18n* to work on this issue. The group is headed by Dr. M. Ohta of Tokyo Institute of Technology in Japan. Internationally, various organizations are working on the local language support, which is called "internationalization and localization." The former provides the framework on which all local languages are supported, and the latter is actual implementation of the local language support. To have a unified view on the localization, the localization profile for each country or region is necessary. The *International Organization for Standardization* (ISO) is working on the internationalization through its technical committee JTC1, and subcommittees. UniForum and X/Open are working together on the subject. Other organizations such as IEEE and OSF are also working on various issues through their working groups.

Among the Asian countries, Japan and the Republic of Korea have fairly well developed local language support in their network environments. Chinese language support is being developed by various countries and regions, and their efforts need to be harmonized. Many other languages are supported in the Internet and other networks, and their usages vary among countries.

We need to work at two levels; one at global and regional levels on internationalization, and the other at an individual country or a language region on localization. We need to come up with the internationalized network software. The localized network software may be handled through local clearinghouses.

Universal code systems such as Unicode and ISO 10640 may have major impact to the internationalization and localization. [5] These systems offer a generic framework which makes localization to specific language environments much easier and consistent. On the other hand, the harmonized development of such a code system is not easy at all.

Link coordination

Most of the international leased lines in the Asia-Pacific region go to the USA. The leased lines to Europe and within the Asia-Pacific region are not popular. There are several reasons for this:

- Intercontinental leased lines cost marginally more than intracontinental leased lines.
- All countries primarily communicate with the USA.
- One fat pipe makes more sense than multiple thin pipes.

By connecting to the same place in the USA or in the region with a fat pipe to the USA, we can solve the problem of intracontinental exchanges among the Asia-Pacific countries. [3] The *Global Internet Exchange* (GIX) is the proposed model. [1, 9] The European group implemented the GIX on the east coast of the US, and is working on its extension to allow a distributed GIX system by supporting multiple GIXs. For the Asia-Pacific community, it is natural if the Asia-Pacific GIX is located in the west coast of the US or in the Pacific region. For the distributed GIX, the problem is who will support the link between GIXs as we need fat pipes for the link. The cost for the Atlantic link is shared between the US and Europe, and a similar scheme may work for the Pacific. For future link coordination, the emerging commercial operations should be taken into account as they may play the major role. CAREN is one of the regional networks which provides Intra-continental links. It links Japan, Korea and Taiwan with the addition of China soon.

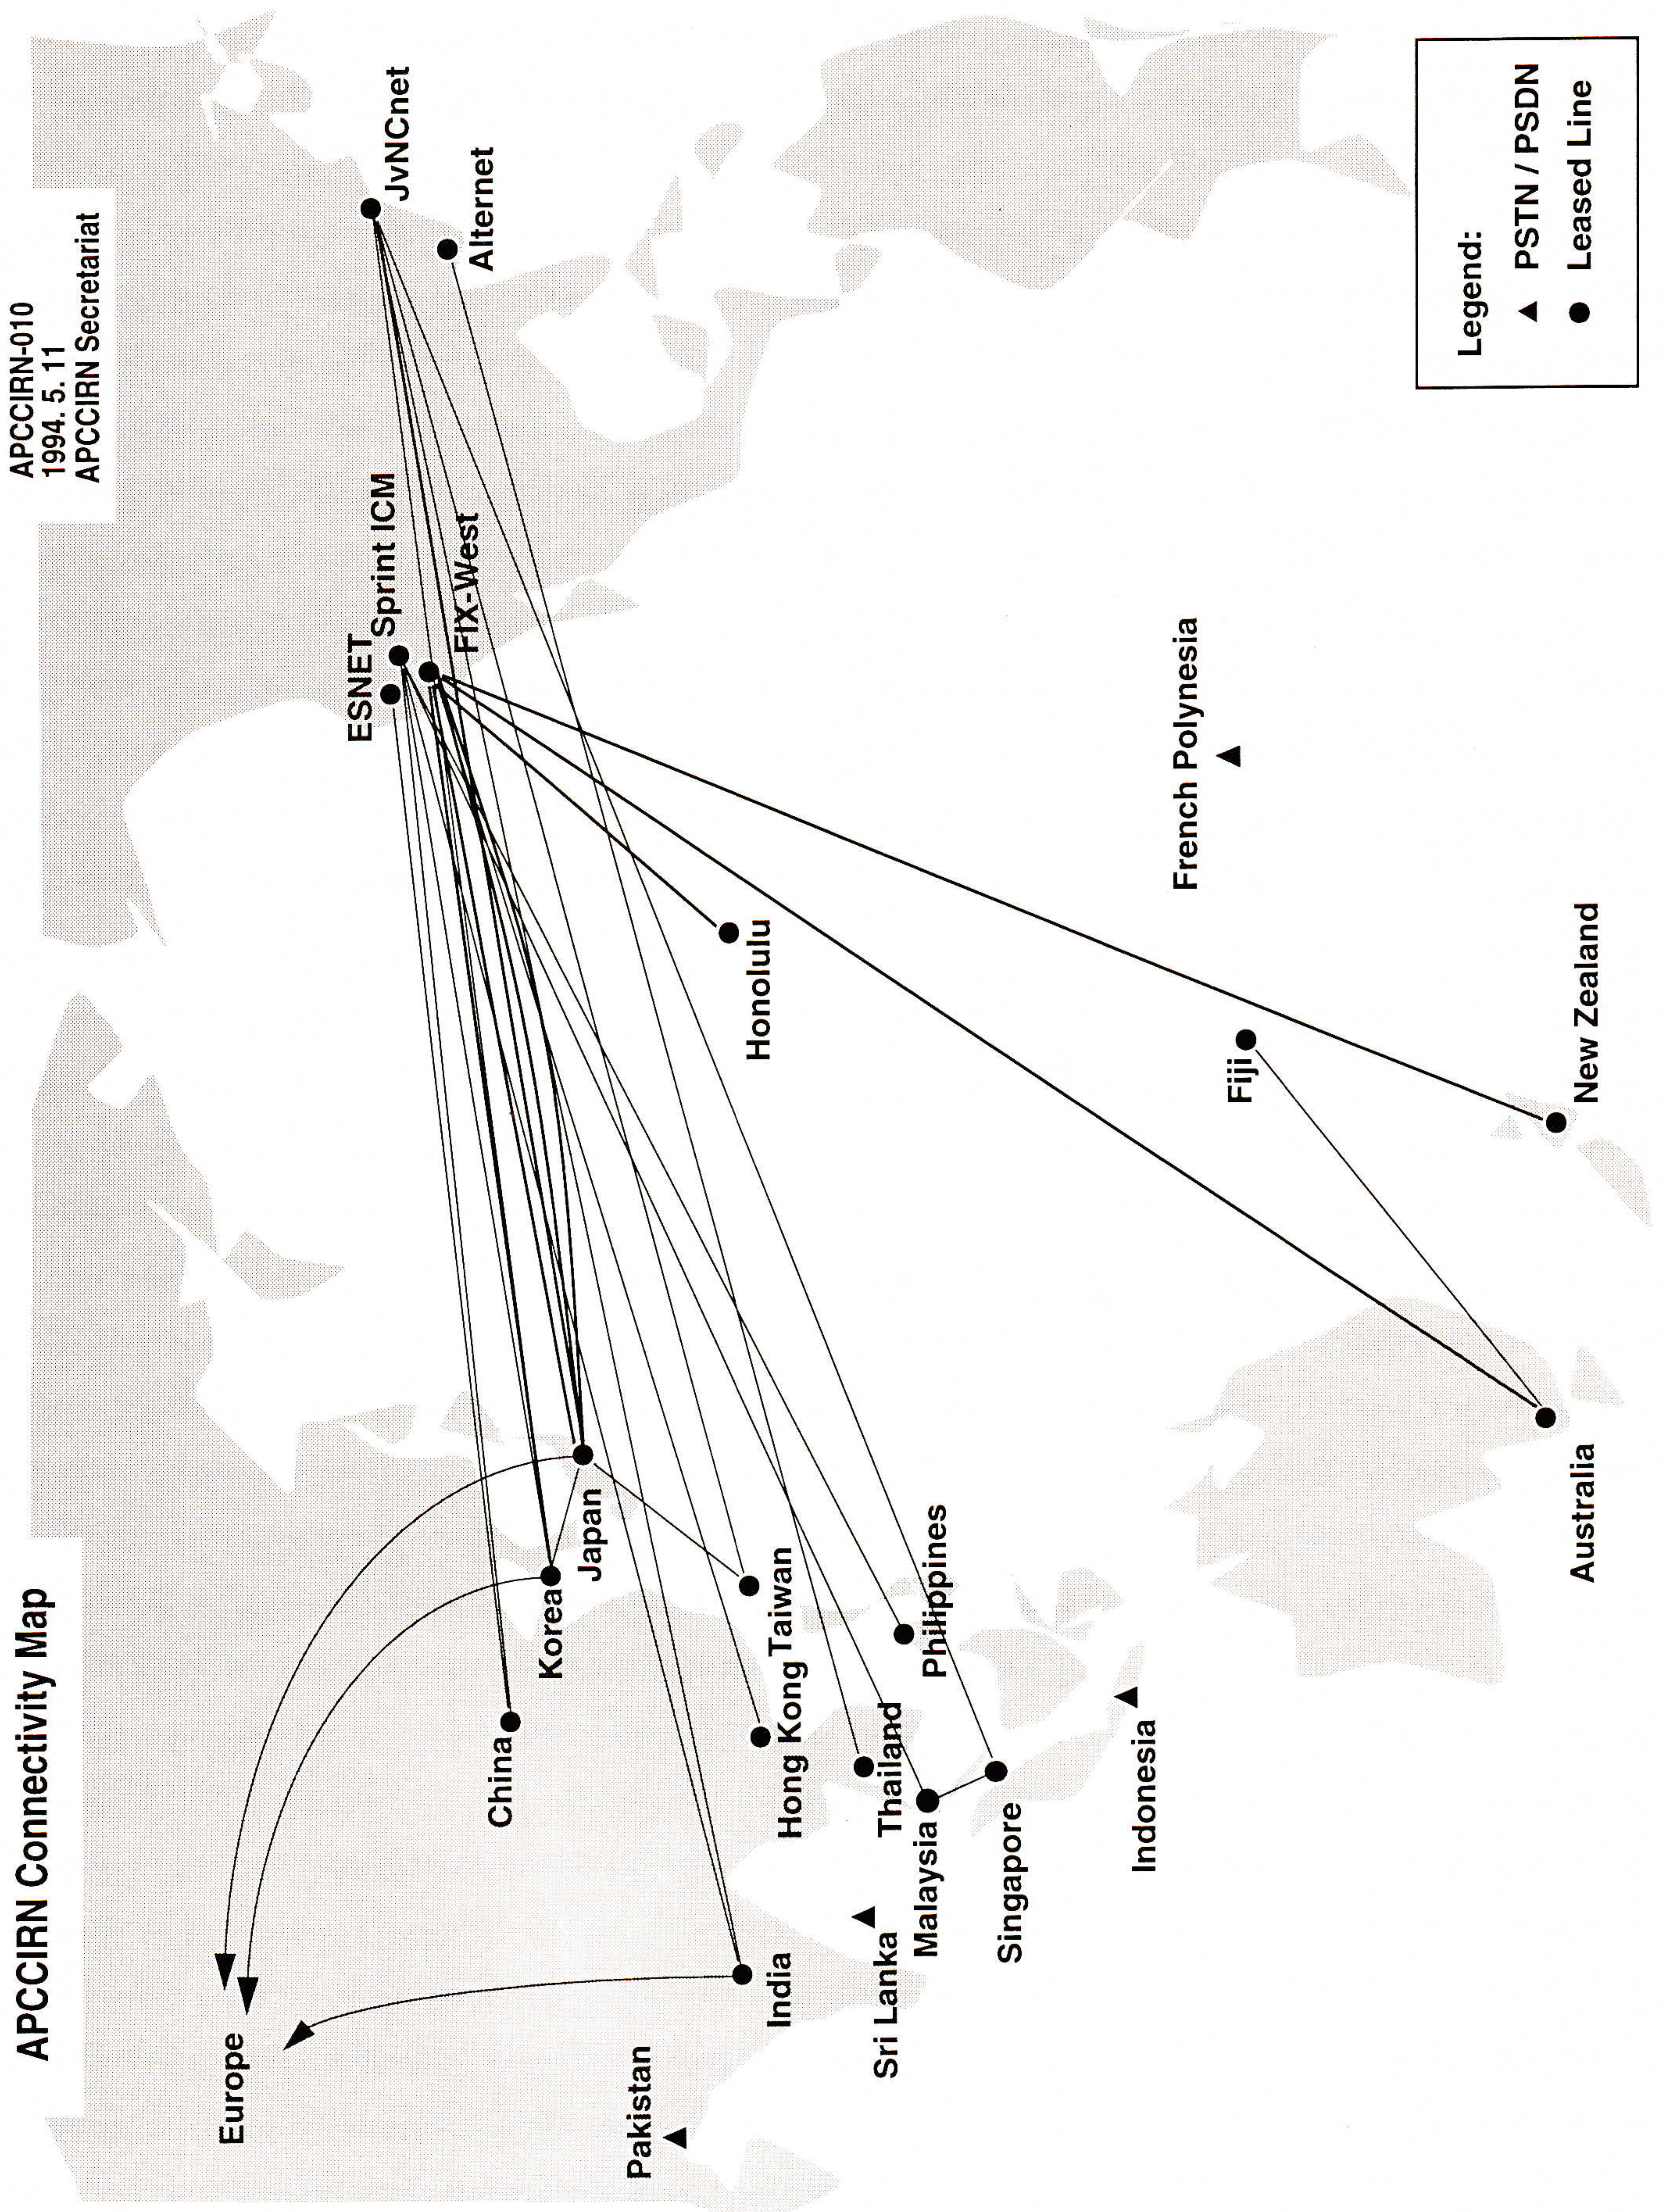


Figure 1: APCCIRN Network Links

continued on next page

Internetworking in Asia (*continued*)

Networked information

Networked information is an emerging major issue to the Asia-Pacific community. As networking becomes more common, information access becomes more important. In addition to general issues of the global networking community, we need to address the issue of the local language support. Most of the information materials in the region are in local languages, i.e., non-English. This causes problems on information tools, translation, and other issues on natural languages.

The *Pacific Neighborhood Consortium* headed by Prof. C. Hardyck of the University of California is addressing the issue of information sharing among the Pacific Rim countries. [7] The consortium was established in early 1990s, and meets annually.

Concluding remarks

The Asia-Pacific networking community may be able to contribute to globalization of the Internet by addressing some of the above issues and others. The globalization of the Internet is very important for the development of good global information infrastructure. As part of the global information infrastructure, an Asian Information Infrastructure based on the Asian Information Highway would be the issue our community should address in the coming years.

References

- [1] G. Almes, et al, "Global Internet Exchange (GIX)," Working Paper, IEPG, 1992.
- [2] APCCIRN, *Gopher* `apccirn@krnic.net`.
- [3] K. Chon, et al(ed.), Proceedings of the Pacific Computer Communications Symposium, 1985.
- [4] K. Neggers, "CCIRN," Proceedings of INET '92, 1992.
- [5] International Organization for Standardization (ISO), "Multi-Octet Code System," IS 10646, 1993.
- [6] L. H. Landweber, "International Connectivity," Version 10, February 1994.
- [7] C. Hardyck, "Pacific Neighborhood Consortium," 1993.
- [8] UNESCO, "Academic and Research Networks in the Asia and Pacific Region," 1992.
- [9] Stockman, B., "Global Connectivity: The Global Internet Exchange (GIX)" *ConneXions*, Volume 7, No. 11, November 1993.
- [10] Conrad, D., "The Asia Pacific Network Information Center: Present and Future," *ConneXions*, Volume 8, No. 7, July 1994.
- [11] *ConneXions*, Volume 7, No. 5, May 1993, "Special Issue: Focus on Europe."
- [12] Guy Almes, Peter Ford, and Peter Lothberg, "Proposal for Global Internet Connectivity," June 1992.
- [13] Malamud, Carl, "The Internet Explorer in Japan," *ConneXions*, Volume 6, No. 5, May 1992.

KILNAM CHON holds Ph.D from the University of California, Los Angeles. He worked at Jet Propulsion Laboratory and Rockwell International in USA on networking related projects before he went to Korea in 1979. He is a professor at Korea Advanced Institute of Science and Technology, and his research interests include networking and distributed processing. He chairs the Asia-Pacific Coordinating Committee on International Research Networking (APCCIRN), and co-chairs CCIRN. He can be reached as `chon@cosmos.kaist.ac.kr`

Construction of a Distributed Software Development Environment

by Ryoichi Hosoya and Shunichi Fukuyama, NTT

Introduction

Nippon Telephone and Telegraph (NTT) has been developing a distributed software development environment based on UNIX workstations and computer networks (LANs and WANs) since 1988, in cooperation with major software development organizations. This environment is intended to improve the productivity of software development, which is becoming increasingly important as services become varied and more advanced [1].

Objectives

The objectives of this activity, called CAE (*Computer Aided Engineering*), are:

- To build a nation-wide infrastructure connecting not only NTT and its affiliates, but also other companies engaged in joint development or hired as subcontractors.
- To provide an environment for distributed software development including a database to promote reuse of software modules.

Development of the backbone network, called *CAE-NET*, began in 1989, led by the NTT Software Laboratories. In the prototype phase, which ended in 1990, the network hierarchy, address allocation, network security, and routing design were established.

In particular, through the security control system, the content of communications can be very explicitly restricted by the gateway processors, allowing use of this network by subcontractors.

In April 1991, network operation were transferred to the NTT Telecommunications Network Sector (now named the Long-Distance Communications Sector) to expand the service area. Nationwide coverage began in October 1991, with over 10,000 nodes from Hokkaido to Kyushu.

This article describes some issues and approaches to improving the software development environment, its implementation concept, present status, and some results from the CAE-NET project.

Benefits

The advantages and disadvantages of the conventional centralized environment using mainframes and of the distributed software development environment using workstations and computer networks are obvious, as has already been shown [1]. Use of workstations offers:

- Higher cost effectiveness than mainframes and typically the optimum equipment investment
- Better performance for faster response, as seen in RISC processors, and
- Advanced man-machine interfacing through high-resolution displays and windows interface;

and the network infrastructure (LAN/WAN) provides:

- High-speed information transmission
- Easy sharing of high-speed printers
- Simplified distributive software development by several persons
- High-speed remote file transmission

A Distributed Software Environment (*continued*)

- Circulation of design information and tools between organizations or remote sites,
- Collaborative development between regions, and
- High-speed transmission of trouble reports and repair information.

Due to recent progress in downsizing, open architecture, and distributed systems, the distributed software development environment is more and more widely used [4].

In constructing the distributed software development environment, the following needs arose:

- *Constructing networks within and between organizations:* The main purposes of computer networks are promotion of information circulation and sharing of computer resources. For these purposes, it is necessary to ensure communication between organizations and to provide security to prevent outside intrusion. It is also necessary to have a maintenance system to handle and control the daily maintenance of network nodes and to deal with problems.
- *Standardization of network architecture among all organizations:* If each organization used its own workstations and network architecture, software circulation on workstations and communication between organizations would be far too difficult. Thus it is necessary to prepare common construction guidelines among all organizations in order to standardize platform and communication protocols. Common manuals are also needed to define network address allocation rules, security control rules, and each organization's responsibilities. With this in mind, a working group (software CAE-WG) from the Software Laboratories with representatives from each NTT software organization has been organized to promote the preparation of common guidelines and the improvement of the overall infrastructure.

Construction guidelines

It is necessary to prepare common guidelines beyond the basic infrastructure regarding the specifications for workstations and gateway processors in order to circulate software easily and to communicate between organizations. In these common guidelines, a basic definition must be included that describes the application program interface to ensure tool portability and the communication interface to provide node-to-node communication, as well as implementation definition to define the detailed version number. In establishing these guidelines, we adopted a policy based on industry standards in order to make it easy to use well-proven methodologies and tools available, and to expand the range of product choice.

For the applications interface, we adopted UNIX, which is established as the industry standard, because of its great ease of operation and variety of functions relating to software development.

For the internal Japanese code, we adopted EUC, which is considered the international JIS code for UNIX, and for the window system, we chose The X Window System, used widely as public domain software.

For communications, we adopted Ethernet at the physical layer, and for site communications, we used a Super Digital telecommunication line. We adopted TCP/IP as the communication protocol. Table 1, on the next page, outlines the software development environment construction guidelines.

Objectives	Define Level		Basic Provisions	Actual Provisions	User Object					
	Items to be Defined				①	②	③	④	⑤	
1. To Obtain Tool Portability (AP Interface)	(1) OS	UNIX		XPG3 (Include POSIX)						
	(2) Programming Language	C, Ada		C (ANSI X3J11/88-159)						
	(3) Japanese	EUC		Ada (ISO 8652-1987)						
	(4) Database	SQL		UJIS	○	○	○	-	-	
	(5) Windows	X-Windows		ISO SQL						
	(6) Communication	TCP/IP Protocol (TELNET, FTP, NFS, RPC)		X11R3/4						
2. To Guarantee Node Communication (Communication Interface)	(1) Protocol			DARPA RFC						
	(a) Layer - 1~2	CSMA/CD		IEEE 802.3						
	LAN	(b) Layer - 3~4	TCP/IP Protocol (IP, TCP, UDP, etc.)		MAC Address (ROM)					
		(c) Layer - 5~7	TCP/IP Protocol (TELNET, FTP, NFS, RPC)		DARPA RFC					
		(d) Transmission Code	Alpha numeric		SRI-NIC IP Address (B class)	○	○	○	○	○
			Kanji		DARPA RFC					
	WAN	(1) Protocol			Common resource allocation rule					
		(a) Layer - 1~2	Private/Public		JIS-X0201-1976					
		(b) Layer - 3~4	TCP/IP (IP,TCP, UDC etc.)		JIS-X0208-1983					
		(c) Layer - 5~7	Address by domain method		HDLC/X21/X25					
	(d) Transmission Code	Alpha numeric		DARPA RFC						
		Kanji		DARPA RFC						
				JIS-X0201-1976						
				JIS-X0208-1983						

User: ① System construction, ② Tool implementer, Object: ③ WS, ④ Target machine, ⑤ GW/Line GW: Gateway

Table 1: An example of Distributed Software Development Construction Guidelines

Manager's Manual

A network manager's manual is essential to ensure security and smooth operation of the network. Computer networks based on TCP/IP in particular need a manager's manual, because such networks assume generally that users have no malice, and communication control is essentially a gentleman's agreement. For this purpose, we have prepared a manager's manual that includes:

- A clear definition of the hierarchical network architecture, consisting of an organizational network and a backbone network,
- The establishment of security control in order to communicate with subcontractors [3], and
- The establishment of a method for preventing errors when setting the routing information at development sites.

Architecture

CAE-NET is a corporate-level national network which connects the regional development sites for software development, maintenance and operation by high-speed telecommunication lines. The network:

- Establishes access points (APs) and connects regional software development organizations,
- Uses Super Digital lines for the backbone lines connecting APs
- Establishes a software database for circulating software development information.

The functions of CAE-NET related to software development are:

- (a) Communication functions
- (b) System manager support functions
- (c) Distributed development functions and support
- (d) Software database functions.

UNIX offers (a),(b), and high-speed file transmission and resource sharing (included in (c), above). The remote testing and development management tools in (c) are prepared by each user.

A Distributed Software Environment (*continued*)

The software database has been established to support the mechanism for sharing and reuse described in the introduction, and offers the corporate-level basic information for software development from the CAE-NET support center.

It is necessary to establish a registration and retrieval mechanism for each organization and to register useful information for common use in order to promote information circulation.

Backbone status

The basic CAE-NET backbone network from Hokkaido to Kyushu was completed in October 1991, and is still growing, as shown in Figure 1. More division networks are being added, mainly at the Telecommunications Software Headquarters and Laboratories. The number of network nodes have exceeded 10,000 and there are more than 500 segments this year as shown in Figure 2. The scale of the network will expand further as the Information Systems Headquarters and other offices begin using it.




Figure 1: Construction status of CAE-NET

Division network status

The Telecommunications Software Headquarters completed construction of the distributed software development environment when they moved to the new Gotanda office. They have developed the INSTEP (*Integrated Support System for ESS Software Production*) system that includes (a) a file-marking support subsystem, (b) a documentation production and management subsystem, and (c) a test support subsystem.

All regional centers have been connected through the computer network and are enjoying its benefits, such as exchange of design information and remote testing [2].

With the trend towards multiple vendors, represented by the *Multi-vendor Integration Architecture* (MIA), the Information Systems Headquarters started to build a distributed software environment using workstations and the CAE-NET in April 1992.

The Laboratories, mainly the Information Communication Network Laboratories, the Communication Switching Laboratories, and the Software Laboratories, have already completed connection to the CAE-NET backbone.

Effectiveness

The engineers do not feel the distance between them when using the computer network and every workstation is connected to every other. For example, they can access a switching system in Tokyo from Hokkaido, or compile software separately on the workstations installed at each location [5].

The actual effects that have begun to appear include (a) abolition of on-site working, (b) reduced number of business trips, (c) reduced file transportation time, and (d) easy, rapid communication and information sharing.




Figure 2: Increase of Segments and Nodes in CAE-NET

Conclusion

This article has described a distributed software development environment that can use commercial off-the-shelf tools because it is based on industry standards such as UNIX and TCP/IP, and whose security control makes it applicable to software development with subcontractors. As network software development progresses, improving software productivity becomes increasingly important. It is necessary to expand the use of the development environment, as well as improve the tools necessary for a switch to multiple vendor and groupware applications. Moreover, it is important to train people on use of the underlying tools and methodologies based on them.

References

- [1] S. Fukuyama, T. Yoneda, H. Asami, S. Katsuragi, "A Concept for a Distributed Software Development Environment," *NTT REVIEW*, Volume 3, No. 1, pp 81–91, 1991.
- [2] H. Nagano, K. Yamaguchi, "Construction of a Distributed Software Development Support Environment for ESS," *NTT R&D*, Volume 39, No. 10, pp 1375–1384, 1990.
- [3] H. Asami, K. Tanaka, S. Fukuyama, "Security Control for a Distributed Software Development Environment," *NTT REVIEW*, Volume 5, No. 3, pp 77–81, 1993.
- [4] M. Aoyama, "Distributed Development Environment: A New Paradigm of Software Development Environment," *Journal of Information Processing of Japan*, Volume 33, No. 1, pp 2–13, 1992.
- [5] M. Suzuki, H. Azuma, "Improvement of User Jobs through Computer Network," *Proceedings of 11th Symposium Quality Control for Software Production*, Union of Japanese Scientists and Engineers, pp 171–178, 1991.

RYOICHI HOSOYA holds an B.E., M.E. from Tokyo University in Japan. Since 1968 he has worked for NTT on a variety of software engineering related projects including development of operating systems and compilers of FORTRAN, COBOL and Ada etc., and various projects concerned with distributed software development environments, internet and multimedia. He is a trustee of Internet Association of Japan. He can be reached via e-mail as: hosoya@sokai.s1.cae.ntt.jp

SHUN-ICHI FUKUYAMA holds an B.E. from Osaka University in Japan. Since 1968 he has worked for NTT on a variety of software engineering related projects including development of operating systems and compilers of APL, FORTRAN and Ada etc., and various projects concerned with distributed software development environments and software process assessment.

The UNCL Project

by Masahiro Taka, UNCL

Introduction

The *Ultra-high Speed Network and Computer Technology Laboratories* (UNCL) were founded on March 30, 1994, in Tokyo. The laboratories are funded mainly by the Japan Key Technology Center, and partially by NEC, Fujitsu and Hitachi. The aim of the laboratories is to carry out the research and development on an integrated system of ultra-high speed networks and computers.

Background

LANs and WANs are being developed with transmission rates of giga bits per second. Furthermore, the performance of supercomputers has been improving with the evolution of parallel processing architecture and LSI devices. High speed workstations have also been evolving year by year. Thus a requirement is arising for supercomputers or high-speed workstations to be interconnected via high speed networks and operated as a single information processing system. However, the information transport speed in current distributed systems is some 100Mbps within the systems, some 10Mbps in intra-offices and some 1Mbps in inter-offices. In order to develop an information processing system operating at higher speeds, new research and development must be carried out to meet the requirement.

The target of this new R&D is to provide a distributed processing system operating at 2.4Gbps by effectively integrating ultra-high speed networks and computers. This ultra-high speed information infrastructure will allow us to realize various applications of natural environment simulations in real time, and interactive visualization of scientific data computations.

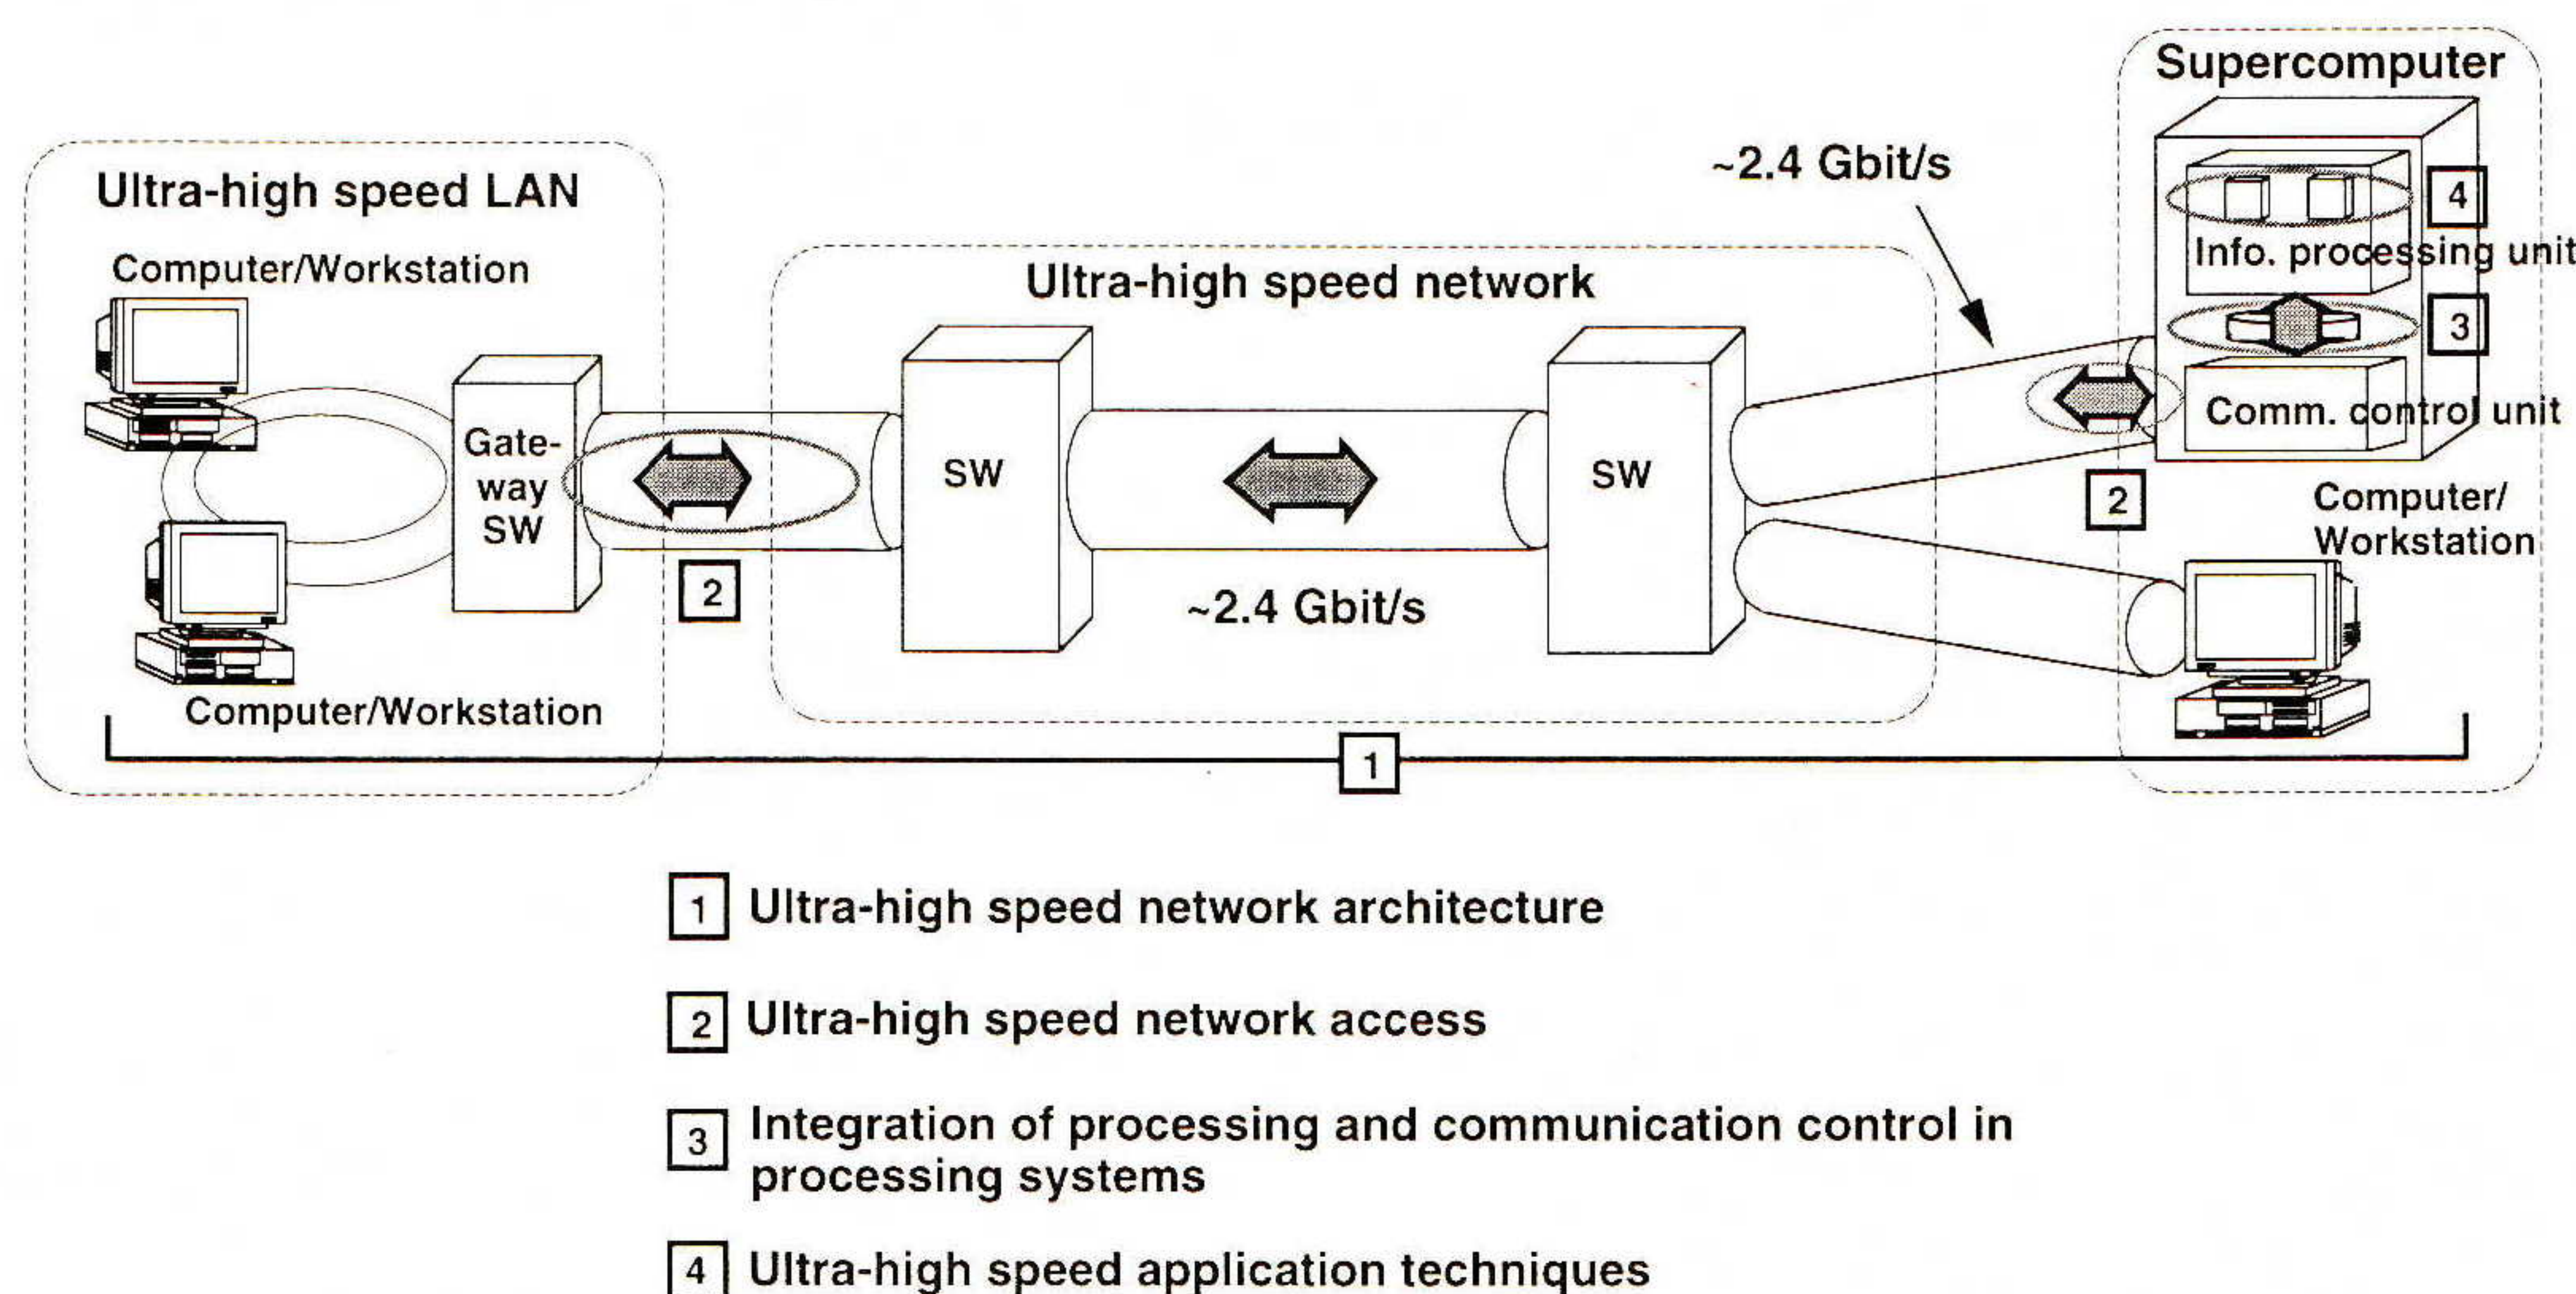


Figure 1: R&D Areas for UNCL

Research areas

The R&D areas and subjects for the project are illustrated in Figure 1. Four research subjects are identified to progress the R&D efficiently:

- *Ultra-high speed network architecture*: Ultra-high speed networking and communication protocols
- *Ultra-high speed network access*: Architecture and protocol processing of ultra-high speed access systems such as gateway systems
- *Integration of processing and communication control in information processing systems*: Input and output channel control and multimedia communication protocols in the processing systems

- *Application techniques*: Highly efficient image signal generation and transport via ultra-high speed networks

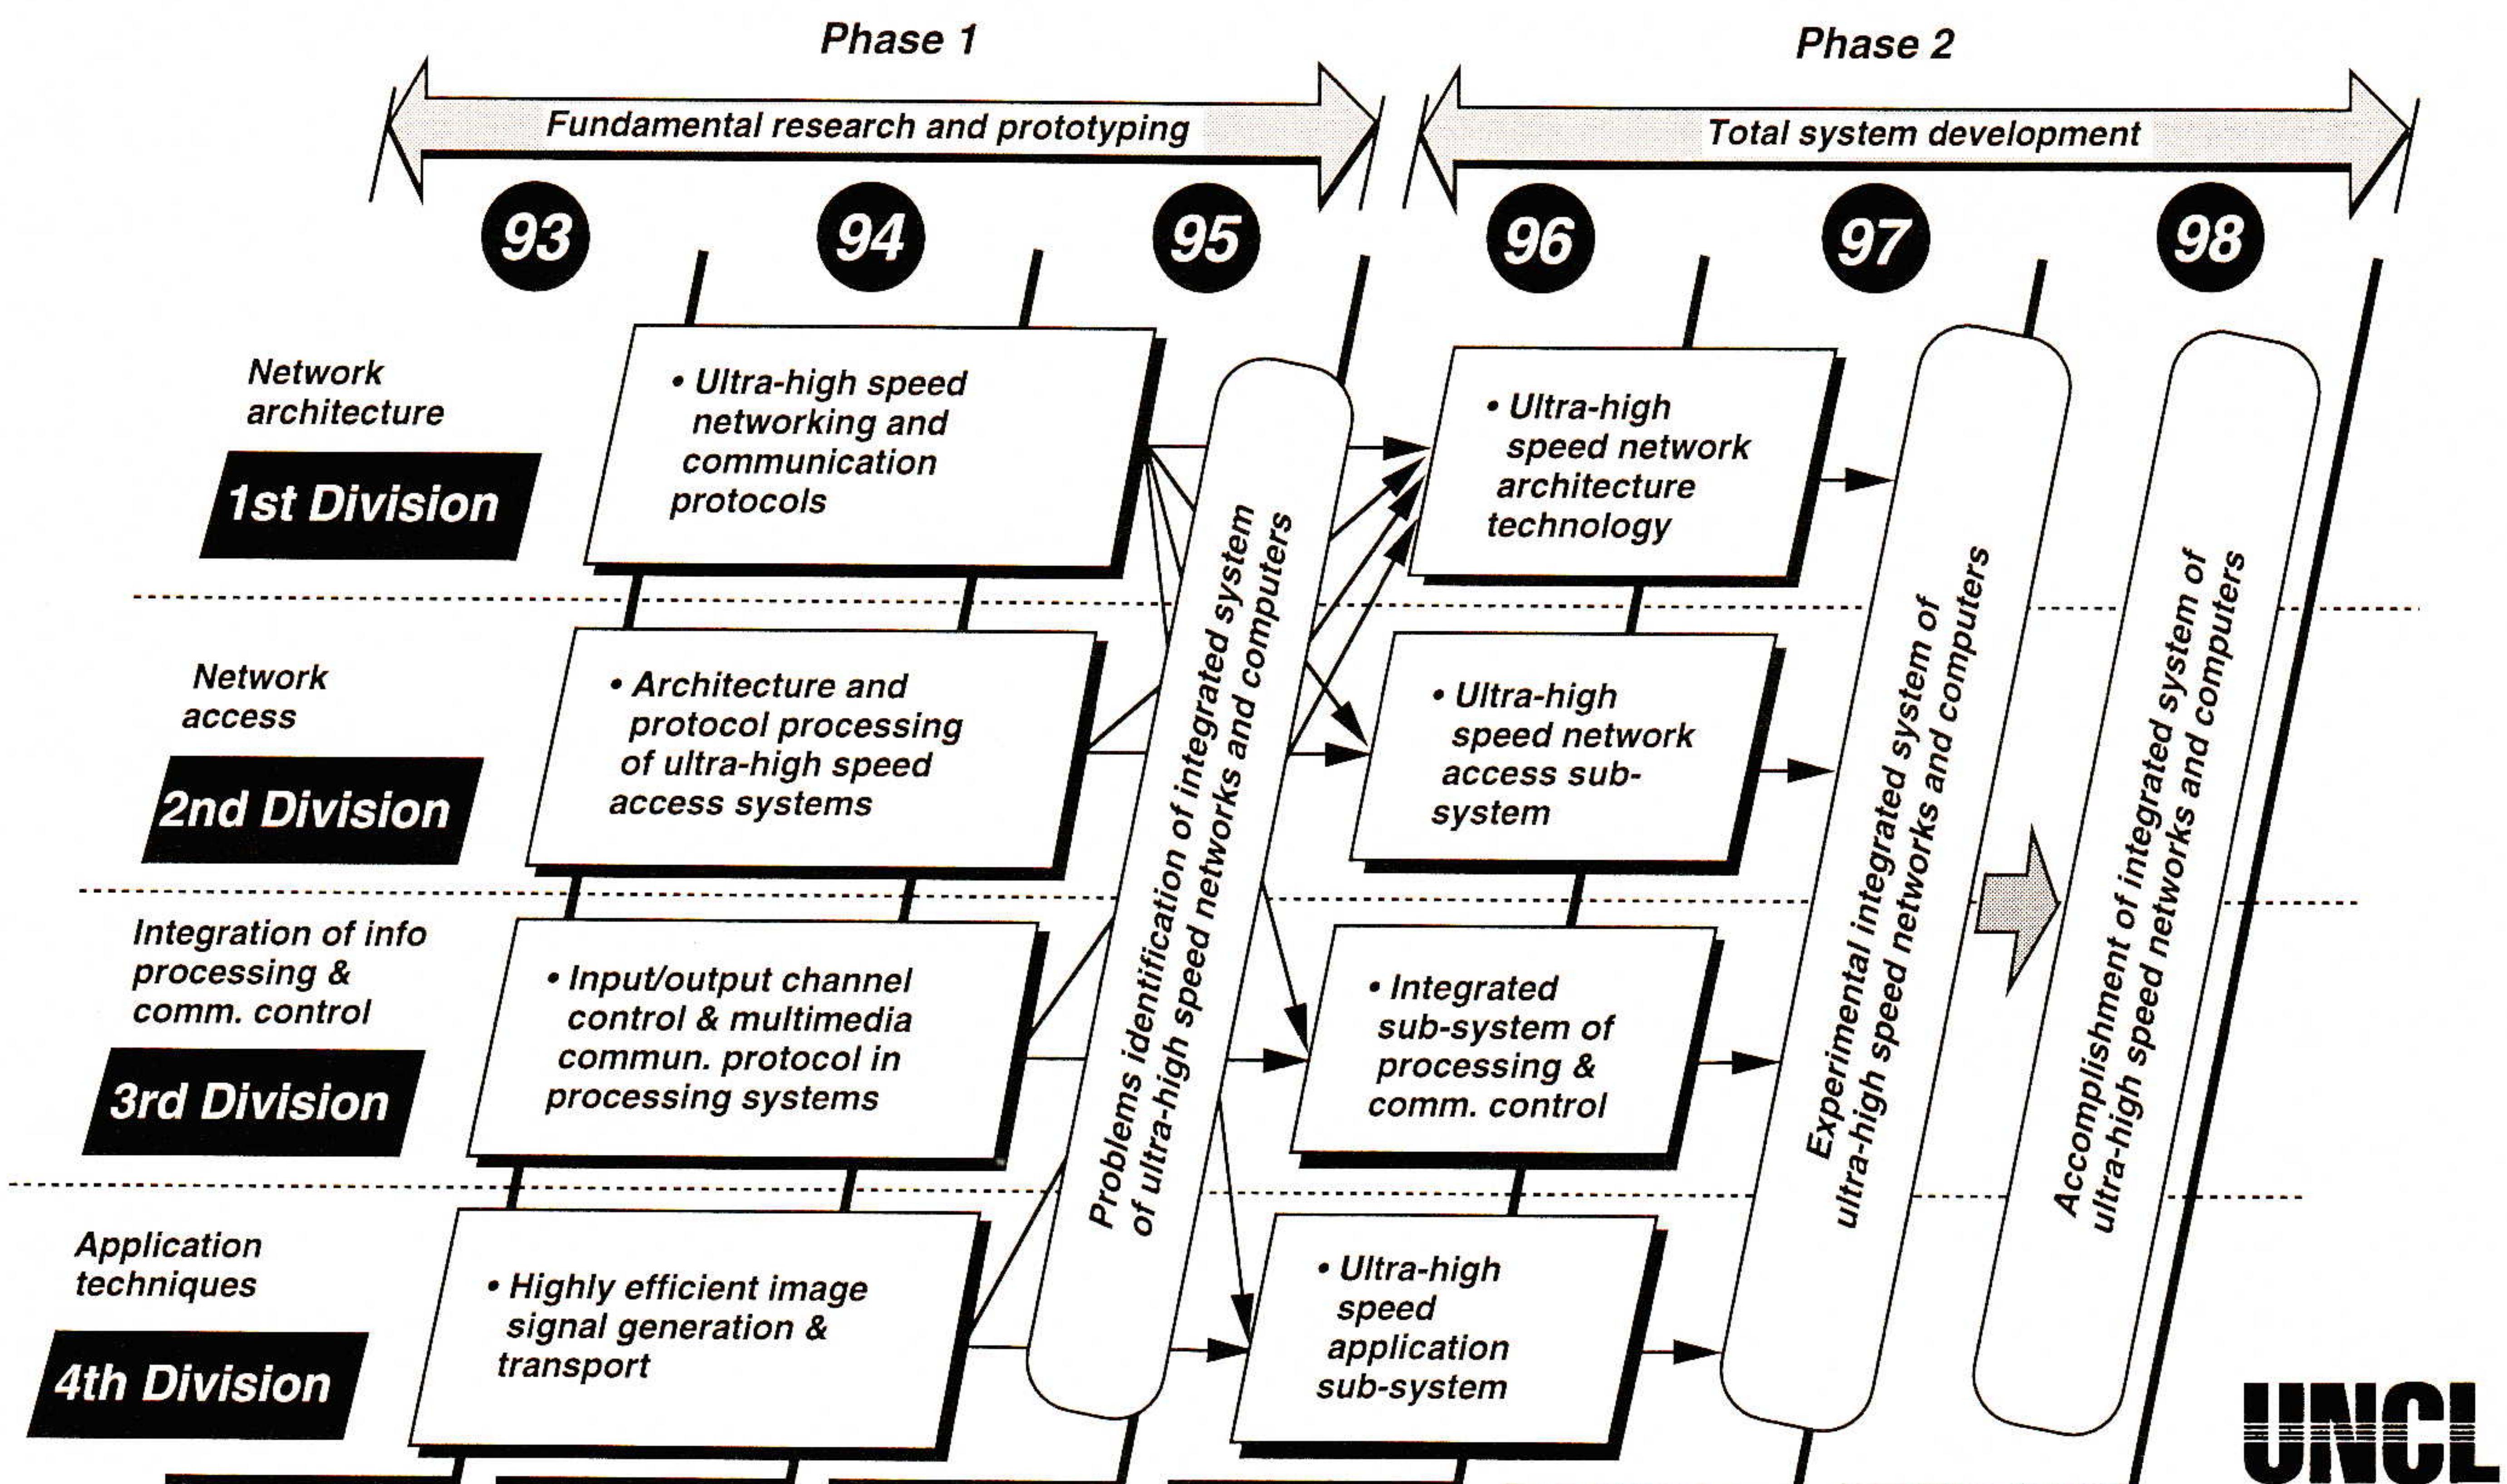


Figure 2: Time schedule for UNCL

Timeframe

The time schedule for the research and development on the integrated system of ultra-high speed networks and computers is shown in Figure 2. The research will be carried out for five years from 1994 to 1998. Phase 1 (1994 to 1995) will be devoted for basic research to explore technical problems toward the development of the integrated system at an ultra-high speed of 2.4Gbps. A laboratory test is planned at the end of Phase 1 to evaluate a prototype system and analyze the performance of the system. Based on the test results, further problems to be solved in Phase 2 (from 1996 to 1998) will be identified. Research will continue to establish the technology for the integrated system of ultra-high speed networks and computers. In 1998, a field test will be carried out to verify the developed system and show the feasibility of the practical deployment.

More information

UNCL
 Toranomom Daini Waiko Bldg. 7F
 5-2-6 Toranomom
 Minato-ku
 Tokyo, 105
 JAPAN
 Telephone: +81-3-3578-9361
 Fax: +81-3-3578-8184

MASAHIRO TAKA holds the B.S. and Ph.D. degrees from Osaka University. Since joining NTT Electrical Communication Laboratories in 1971, he had been active in several R&D areas concerning digital multiplexing terminals, speech coding systems and telecommunication quality techniques. He moved to the Ultra-high Speed Networks and Computer Technology Laboratories in March, 1994, where he is the director of the Labs. Dr. Taka is a member of the IEEE, and of the Institute of Electronics, Information, and Communication Engineers in Japan.

The Asia Pacific Network Information Center: Present and Future

by David Conrad, Internet Initiative Japan, Inc.

Introduction

As is perhaps well known, the Internet in the Asia Pacific region is growing at least on par with the rate of growth of the Internet as a whole. Leased lines, previously rare due to the ridiculously high cost of transpacific telecommunications, are now sprouting up in places such as the Philippines, Indonesia, and Guam. Commercial providers have begun operations in Hong Kong and Japan among other locations, and existing connections are getting faster and wider. In an effort to provide some support to this growth, the *Asia Pacific Coordinating Committee for International Research Networks* (APCCIRN) and the *Asia Pacific Engineering Planning Group* (APEPG) have undertaken the creation of a network information center for the Asia and Pacific Rim regions, to be known as the *Asia Pacific Network Information Center* or APNIC.

Goals

In January, 1993, at the first APCCIRN/APEPG meeting, the APNIC experiments were initiated, primarily targeting the interactions between the AP regional NIC and other regional NICs, the AP region national NICs, and the Internet users in the Asia and Pacific Rim regions [1]. Later, at the APCCIRN/APEPG meeting following the INET '93 conference in San Francisco, the APNIC experiments were expanded and a pilot project was conceived and initiated [2]. The APNIC pilot project goals were and are fairly succinct:

- Determine the requirements for a regional NIC and the means to meet those requirements
- Implement a regional IP address allocation strategy in accordance with RFC 1466 [3]
- Provide a testbed for experimentation into network coordination in the Asia Pacific region
- Coordinate with local, national, and regional NICs
- Experiment with tools used to support NIC operations

In order to insure timely results, the APNIC pilot project was chartered to begin operation on Sept. 1, 1993, and end on June 30, 1994.

As a means to meet these project goals, as well as to provide services to the AP region's networking communities, the APNIC pilot project has been coordinating with the IANA, and the regional registries, InterNIC in North America and RIPE-NCC in Europe. This coordination has taken the form of discussions with RIPE-NCC regarding various aspects of running a regional registry and consultation with InterNIC and the IANA on address space allocation issues in the AP region. On April 1, 1994, this coordination, particularly with InterNIC and the IANA led to the APNIC pilot project officially receiving the delegation of the 202.x.x.x and 203.x.x.x address blocks. Since that date, the APNIC pilot project has been assigning IP addresses to organization in the AP region, maintaining the authoritative database for networks in the 202 and 203 blocks, as well as providing IN-ADDR.ARPA domain space for those blocks.

After eight months of operation, and a month of handling address assignments for the Asia and Pacific Rim regions, the APNIC pilot project has grappled with many of the issues involved with providing a regional NIC service and has developed several proposals regarding the establishment of a permanent APNIC.

This article will present the organizational and funding models of the APNIC pilot project and some of the ideas the pilot project staff have on organizational and funding models of the permanent APNIC. Also, some of the lessons learned during the pilot project will be discussed, and finally some concluding thoughts regarding the work towards a regional network information center in the Asia Pacific region will be presented.

APNIC organization

Given the immense size of the Asia and Pacific Rim regions, from the Persian gulf area to the island nations of the South Pacific, and the vast diversity of cultures, religions, and economic situations encompassed in this space, one of the basic assumptions held by the members of the APNIC pilot project staff regarding the establishment of a NIC in the AP region was that the ultimate APNIC would need to be highly distributed. The APNIC pilot project, following this premise, exists primarily as a set of mailing lists on a machine at the University of Tokyo in Japan. The staff mailing list implementing most of the APNIC functions currently consists of 25 people from the countries of Australia, China, Japan, Korea, New Zealand, Singapore, Taiwan, Thailand, and the US. This informal group comes to a rough consensus on requests for information and/or address space assignment and fulfills those requests typically within one to two working days.

In terms of physical existence, again, distributing the responsibilities is the means and the goal of operating the pilot project. The APNIC pilot project currently shares a Sun SparcStation of the Japanese national NIC, JPNIC. On this machine, which has as one of its names `apnic.net`, the pilot project maintains mailing lists, APNIC archives (available via anonymous FTP and *Gopher*), the Asia Pacific network database and *whois* server (currently running the InterNIC's *rwhois* server software), and a DNS server holding a set of CNAMEs to machines which provide the APNIC services. One of these CNAMEs, `ns.apnic.net`, points to a host operated by the Australian national NIC, AUNIC, which maintains the 202 and 203 `IN-ADDR.ARPA` domains. Another CNAME, `www.apnic.net`, points to a machine operated by the Korean national NIC, KRNIC which houses an experimental WWW home page for the Asia and Pacific Rim regions. By distributing the NIC services, the APNIC pilot project has been able to take advantage of talents of personnel at the national NICs and thus does not require extensive expertise to be located centrally at an APNIC facility.

With the experiences gained from the distribution of both the decision making process and the project hardware, the APNIC pilot project staff has proposed that the permanent APNIC should be distributed functionally to various national NICs [4]. This model of organization is similar in concept to that of the InterNIC in the US which is comprised of three commercial companies, AT&T, General Atomics, and Network Solutions, each providing a specific part of the InterNIC's functionality [5]. In the case of APNIC however, the NIC services will be distributed to national NICs instead of commercial companies. This organizational model would establish several independent functional areas which would be assigned to national NICs for some "contracted" period of time, transitioning to another national NIC when the "contract" expires. This rotation of the responsibilities of APNIC is seen as a way of generating technical expertise in the area of providing NIC services in the various countries as well as serving as a way to provide the services required of a regional NIC.

Asia Pacific Network Information Center (*continued*)

Coordination

With such decentralization, one obvious concern is overall coordination. In order to address this concern, the pilot project staff proposes that the services provided by the national NICs should be coordinated by a small organization within APNIC, tentatively named the *APNIC Coordination Center* (APNCC). This coordination center would be ultimately responsible for the proper operation of APNIC and would supervise the transition of APNIC functions from one national NIC to another, providing technical advice and support as necessary. As a further responsibility, the APNCC would provide the services for which no national NIC is willing or able to accept responsibility. Thus, the proposed organizational plan would provide for an APNIC that would exist as a cooperative organization of national NICs, coordinated by the APNCC.

While all the functions that APNIC should provide have not been fully defined, the APNIC pilot project is currently experimenting with this functional delegation scheme as previously mentioned. These experiments have progressed exceedingly well, primarily due to the expertise and talent of the individuals implementing those APNIC services. With the assumption that additional talented people from other national NICs will be able to provide their talents, there is reason to believe the KRNIC proposal will prove to provide a functional organizational model for the ultimate APNIC.

Funding

The APNIC pilot project is funded in its entirety by the Japanese national NIC, JPNIC which has very generously allocated 10% of its operating funds for the duration of the APNIC pilot project [2]. Obviously, since APNIC will be an organization providing support to the entire Asia and Pacific Rim regions, it is neither desirable nor likely this funding situation will continue. As one of the APNIC pilot project's primary goals, the determination of a stable funding model has been under significant discussion. For the APNIC, a stable funding model would be one in which the APNIC can have some level of assurance of continued existence, without the need for the staff of APNIC to spend most of its time scrounging up money.

In the process of researching the issues of a stable funding model, the APNIC pilot project staff has looked at the means by which InterNIC and RIPE-NCC, the other regional Internet registries, obtain their funding. The two registries, while sharing some of the same functionality, have radically differing budgets and funding models.

The InterNIC is primarily funded by a grant from the US National Science Foundation with the three parts of the InterNIC being funded differently [6]. For Registration Services, the NSF provides full funding via a grant of just over US\$ 1,000,000/year for the duration of the 5 year cooperative agreement. The Database and Directory Services award has a total cost of approximately US\$ 2,000,000/year with approximately US\$ 600,000 coming from NSF, and the remaining cost being split approximately equally between cost-sharing and project related income. Information Services totals approximately US\$ 1,300,000/year with NSF providing a declining amount each year, the remainder of the funds being generated by project related income.

RIPE-NCC, the European regional Internet registry, which has a total 1994 budget of ECU 260,000 [7] has given much thought to the issues of funding models for Internet registries (in particular, see the RIPE document, *ripe-084* [8]). Initially, RIPE derived its funding from RARE, but has developed a funding model that relies on "voluntary contributions" from the IP service providers RIPE-NCC supports.

In the event that the IP service providers do not provide enough money for the operation of RIPE-NCC, RARE guarantees sufficient funding for the continued existence of RIPE-NCC.

One aspect shared between both the RIPE and InterNIC funding mechanisms is the existence of a external funding authority which can provide economic assistance. Within the Asia Pacific region, no such organization has been identified. Thus, the APNIC funding model needs to take into account the fact that there is no single organization to which the APNIC can turn for money. In addition, the RIPE funding model implies the existence of a sufficient number of service providers with sufficient cash reserves to make voluntary contributions. Perhaps due to the high cost of connectivity, the AP region does not at this time have the wide variety of Internet service providers, especially in the commercial realm, as exists in Europe. Thus, due to the nature of the AP region, the APNIC pilot project staff has been unable to make a convincing argument for implementing either the InterNIC funding model nor the RIPE funding model.

The APNIC pilot project staff has therefore concentrated on leveraging the APNIC organizational model. The inherent assumption of the APNIC organizational model is that national NICs will be providing APNIC functionality as an adjunct to functionality they must provide for their own constituency. If this assumption is extended to providing some part of the national NICs funding to support the operation of APNIC, the APNIC funding issue may be resolved.

Even though the APNIC pilot project staff expended significant cycles on funding models, no firm decisions has been made at this time. This issue, a significant problem for not only APNIC, but for the other regional registries as well, will need to be discussed in detail at the upcoming APCCIRN/APEPG meeting in Prague before a stable funding model can be assured.

Lessons learned

At this early stage, the APNIC pilot project has been receiving a fair amount of traffic and that traffic has proven quite instructive to the pilot project staff. Prior to the acceptance of delegation of the 202 and 203 address blocks, the APNIC pilot project staff received 1–5 messages a week, mostly for information regarding Internet connectivity in the AP region. Since the delegation, the pilot project staff receives approximately 5 to 10 e-mail messages and 1 to 5 fax messages a day. It is expected in the future, the number of messages the APNIC staff receives will grow proportionally to the growth of the Internet in the AP region. From the experiences gained in handling this level of traffic, the APNIC pilot project staff have learned several lessons.

Perhaps foremost among the lessons learned is that which should be fairly obvious, namely that running a regional NIC requires significant investments in both time and talent. A regional NIC becomes a magnet, both for informational questions such as “my girlfriend lives in Uzbekistan, how can I find her e-mail address?” to technical questions such as how to use variable length subnet masks to make the greatest use of small amounts of address space. The former type of question, while typically not technically challenging, can be time consuming to answer and requires knowledge of how to navigate the Internet to search for possibly well hidden data. The latter type of question typically requires extensive knowledge of the current popular techniques for building and maintaining IP networks, and can also require significant amounts of time to resolve.

Asia Pacific Network Information Center (*continued*)

The APNIC pilot project staff has also learned the functions associated with running an Internet registry such as address allocation can not only be time consuming and require significant knowledge regarding routing techniques and protocols, but also can require the individual responding to the request to appear nearly monomaniacal with respect to not allocating the requested amount of address space if the request is not sufficiently justified. Since allocations of network addresses must be carefully considered, with special emphasis placed on the allocating the appropriate amount of address space in a way that conforms to the requirements of CIDR, the APNIC project staff has had to explain Internet addressing and the global routing table situation many times, sometimes more than once to the same individual. In the AP region, this sort of situation can have added complexity due to language and cultural differences and typically must be handled with some care.

Thus, briefly stated, the primary lesson the APNIC pilot project staff has learned is that running a regional NIC is not for the weak of heart, shallow of mind, or shy of disposition. The pilot project staff has found that, in order to provide a high level of service both to the clients of APNIC, and to the Internet at large, APNIC staff must be willing to spend significant time educating requesters, be thoroughly knowledgeable in a wide variety of networking technologies and techniques, and must be able to say "no," albeit politely, when appropriate.

Conclusions

In terms of the future APNIC, it is believed by the APNIC staff that the organizational model proposed will provide a flexible and effective means to ensure APNIC can operate effectively given the Asia and Pacific Rim regional constraints. The distributed model has been shown to be workable both in the US with the InterNIC, and via experiments conducted within the APNIC pilot project, and should allow APNIC to provide timely and effective services.

With respect to the APNIC funding model, basing the funding mechanisms upon the concept of leveraging the APNIC organizational model should provide a way of ensuring enough funding to make APNIC viable. It is obvious however that additional work is required in this area, perhaps in concert with the other regional registries.

Obviously, for the APNIC pilot project, the most significant event in the future is the conclusion of the pilot project phase on June 30, 1994. The pilot project staff will be presenting a final report on aspects of the project and will be presenting this report to the APCCIRN/APEPG at their meeting following the Prague INET meeting. It is believed that the results of the pilot project should provide enough information, techniques, procedures, and software to enable a smooth transition to the permanent APNIC.

References

- [1] Murai, J., & Hirabaru, M., "A Proposal for APNIC Experiments," <ftp://nic.nm.kr/apccirn/020.apnic.proposal>, January 1993.
- [2] Murai, J., & Conrad, D., "Asia Pacific Network Information Center Pilot Project Proposal," <ftp://apnic.net/apnic/docs/english/apnic-003.txt>, August 1993.
- [3] Gerich, E., "Guidelines for Management of IP Address Space," <ftp://ftp.iij.ad.jp/pub/rfc/rfc1466.txt>, May 1993.

- [4] Park, T., "Proposal for Distributed APNIC Operations," `ftp://apnic.net/apnic/docs/english/apnic-007.txt`, December, 1993.
- [5] US National Science Foundation, "Network Information Services Manager(s) for NSFNET and the NREN, Project Solicitation," National Science Foundation, March 1992.
- [6] Mitchell, D. and Goldstein, S., Private communications September, 1993.
- [7] Karrenberg, Daniel, "RIPE NCC Funding Situation," message to `ripe-list@ripe.net` electronic mailing list, October, 1993.
- [8] Blokzijl, Rob, "RIPE NCC Funding," `ftp://ftp.ripe.net/ripe/docs/ripe-docs/ripe-084.txt`, May, 1993.
- [9] Karrenberg, Daniel, "The RIPE NCC and the Routing Registry for Europe," *ConneXions*, Volume 7, No. 11, November 1993.
- [10] Tony Bates, Jean-Michel Jouanigot, Daniel Karrenberg, Peter Lothberg and Marten Terpstra, "Representation of IP Routing Policies in the RIPE Database," March 1993
- [11] Stockman, B., "Global Connectivity: The Global Internet Exchange (GIX)" *ConneXions*, Volume 7, No. 11, November 1993.
- [12] *ConneXions*, Volume 7, No. 5, May 1993, "Special Issue: Focus on Europe."
- [13] Guy Almes, Peter Ford, and Peter Lothberg, "Proposal for Global Internet Connectivity," June 1992.

DAVID "RANDY" CONRAD holds a B.S. in Computer Science from the University of Maryland, College Park. He has been working in the networking field since 1984 and has been involved with networking in the Asia and Pacific Rim region since 1990, initially working with the PACCOM project at the University of Hawaii. Since 1992 he has worked for Internet Initiative Japan, Inc., a commercial service provider in Japan, as a R&D engineer. In addition, since September, 1993, he has also been the technical contact for the APNIC pilot project and a member of the APNIC pilot project staff. He can be reached as: `davidc@apnic.net`

Coming to *ConneXions*

Now is a great time to renew or start your subscription. Take advantage of the 20% conference discount (only \$120 domestic or \$170 international for one year). Here are some of the articles we are working on for future issues:

"RSVP: A Resource ReSerVation Protocol"
by Bob Braden and Lixia Zhang

"Standards: When is it too much of a good thing?"
by Robert J. Aiken and John S. Cavallini

"Windows Sockets 2.0"
by J. Allard

"Interworking with B-ISDN"
by Reto Beeler

"Cryptography in computer networks"
by William Stallings

"Corporate IP Networks"
by George Abe

IP Over ATM and the
Construction of High-Speed Subnet Backbones

by Mark Laubach, Com21, Inc.

Introduction

IP over *Asynchronous Transfer Mode* (ATM) promises to represent a feasible solution for building gigabit logical IP subnets (LISs). But just how hard is it to really build one? This article gives an overview of the *Internet Engineering Task Force* (IETF) IP over ATM proposed standards and then summarizes the engineering considerations for the construction of the *Bay Area Gigabit Testbed* (BAGNET).

ATM implementation
in review

Yes, ATM is here but you must ignore all the marketing hype and the Washington-centric lobbying for future options about the “National Information Superhighway” in order to see beyond the billboards—the road is unpaved and there are many potholes for the unsuspecting tourist. ATM in the vendor community is still in its infancy and it will take a few decades to sort things out before we see ATM “everywhere” it is promised to be.

The *good* news is that ATM is here in its infancy, it is indeed fun, promising, and very exciting technology. It offers the speed and performance that we need for the future; it promises the ability to handle aggregate traffic streams of differing *Quality Of Service* (QOS), e.g., mixing your real-time teleconference with your background file transfer; and ATM has the potential of providing global, infinitely scalable connectivity from LANs to long haul links, from back-bones to bread-boxes. Yes, there are technology investigations in progress that are trying to bring ATM to your TV set top over CATV—imagine when it reaches the kitchen [1].

ATM will be globally successful only when the following is available across all implementations: signaling, congestion control, traffic management, multimedia support, internet interoperability, and security [1]. Additionally, the socialization of ATM into the international telecommunications infrastructure will take a few decades to complete.

Encapsulation of IP
over ATM

ATM packages data in units of 53 octet quantities called *cells*. These cells could be called “data photons,” as they are the smallest unit of particles (ignoring bytes) in the ATM networking world. Packets then are more like waves, and ATM protocol data units (packets) are layered on top of the cell structure. The *ATM Adaptation Layer* or AAL for short, gives meaning to a given stream on cells which are transmitted on a particular virtual channel. One particular adaptation layer, called AAL5, is of interest to us in packet-land, as AAL5 layers a 65K protocol data unit on top of the cell stream. See Figure 1.

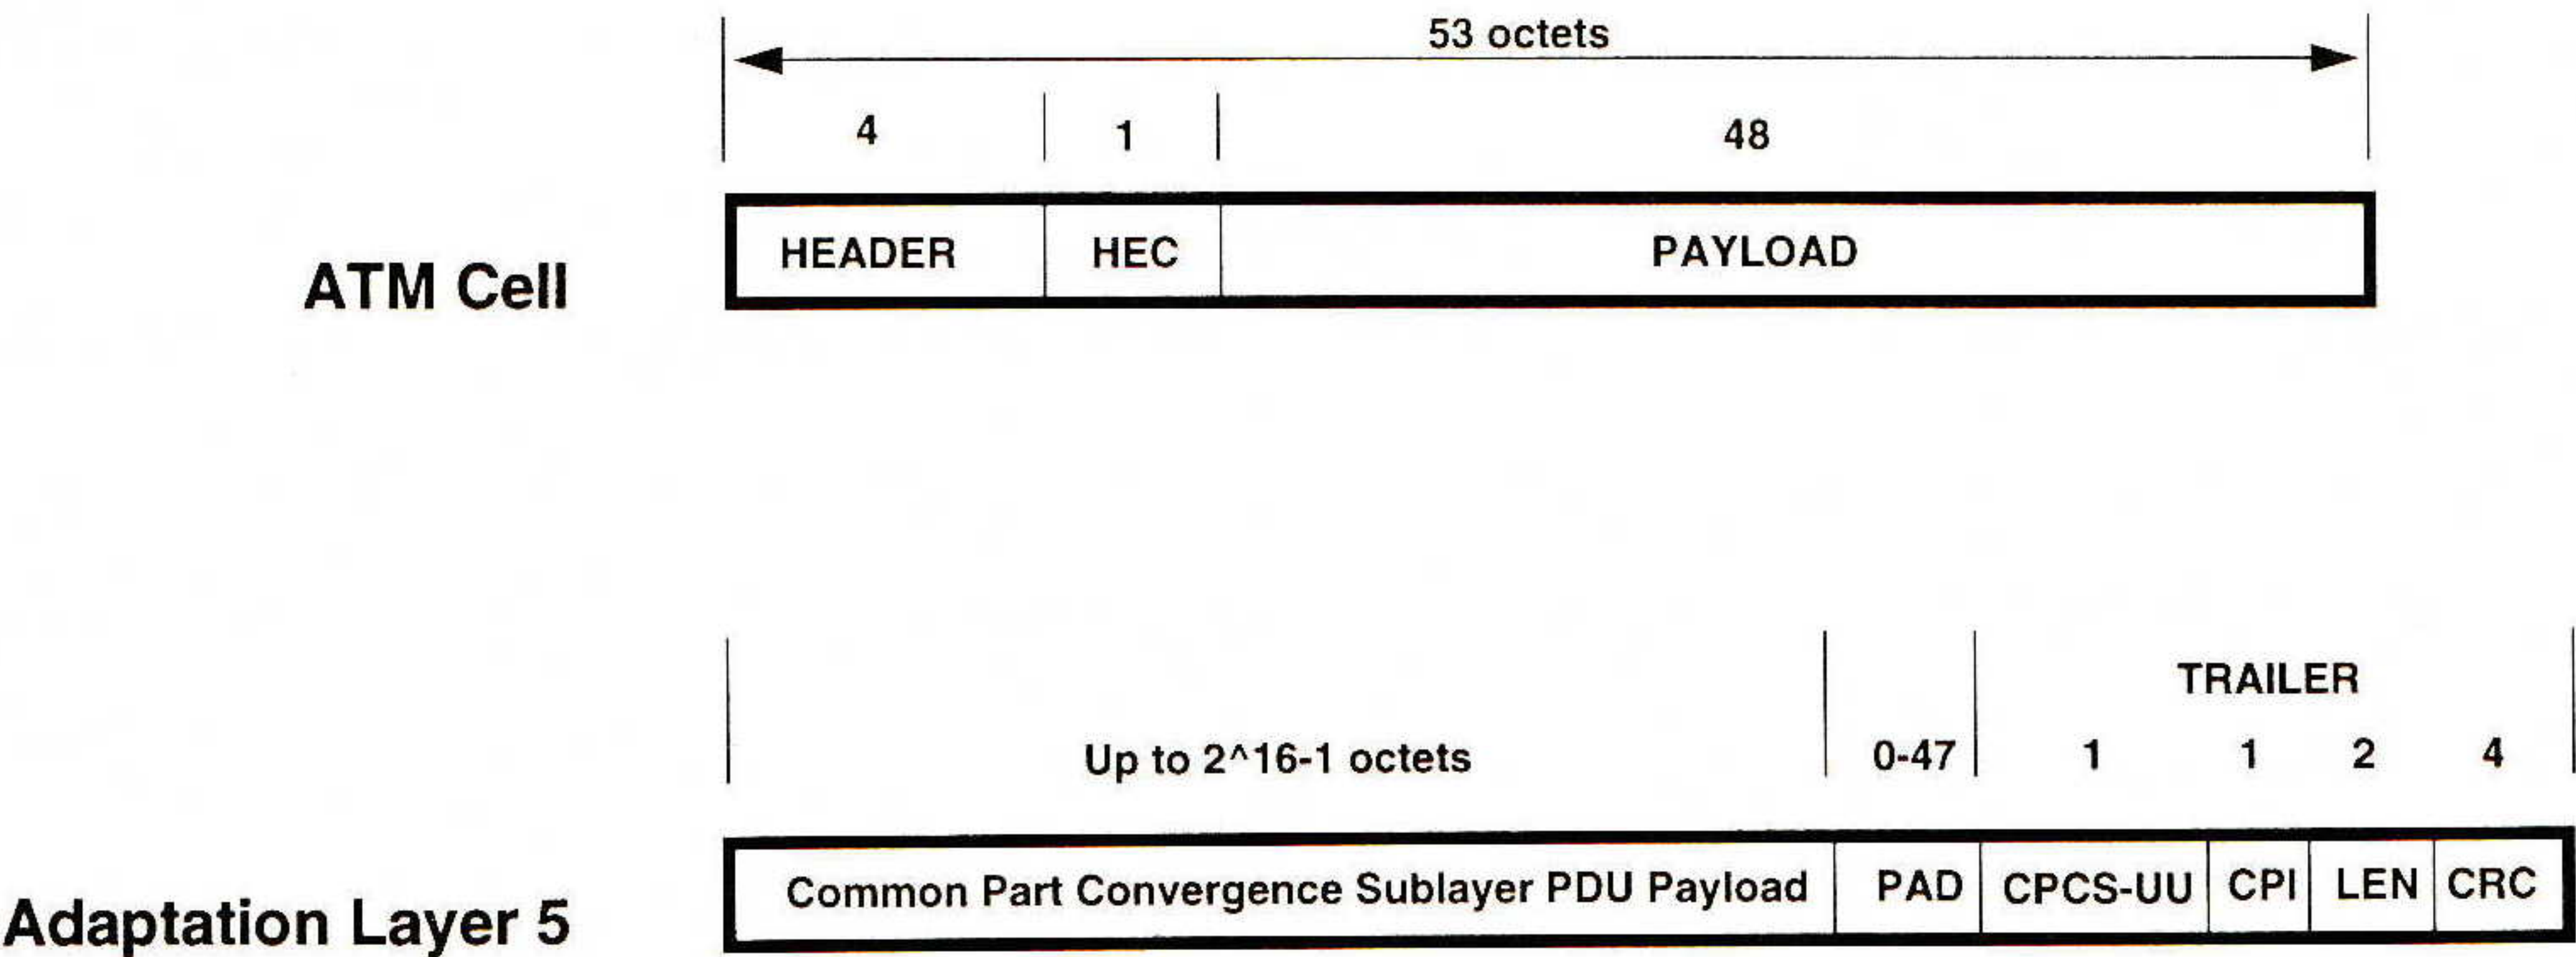


Figure 1: ATM Cells and ATM Adaptation Layer 5

RFC 1483 defines two encapsulations of IP packets into AAL5 protocol data units [3]. The first being LLC/SNAP encapsulation, the second null encapsulation. RFC 1577 has defined LLC/SNAP as being the default for IP over ATM subnets [4]. LLC/SNAP encapsulation is shown in Figure 2.

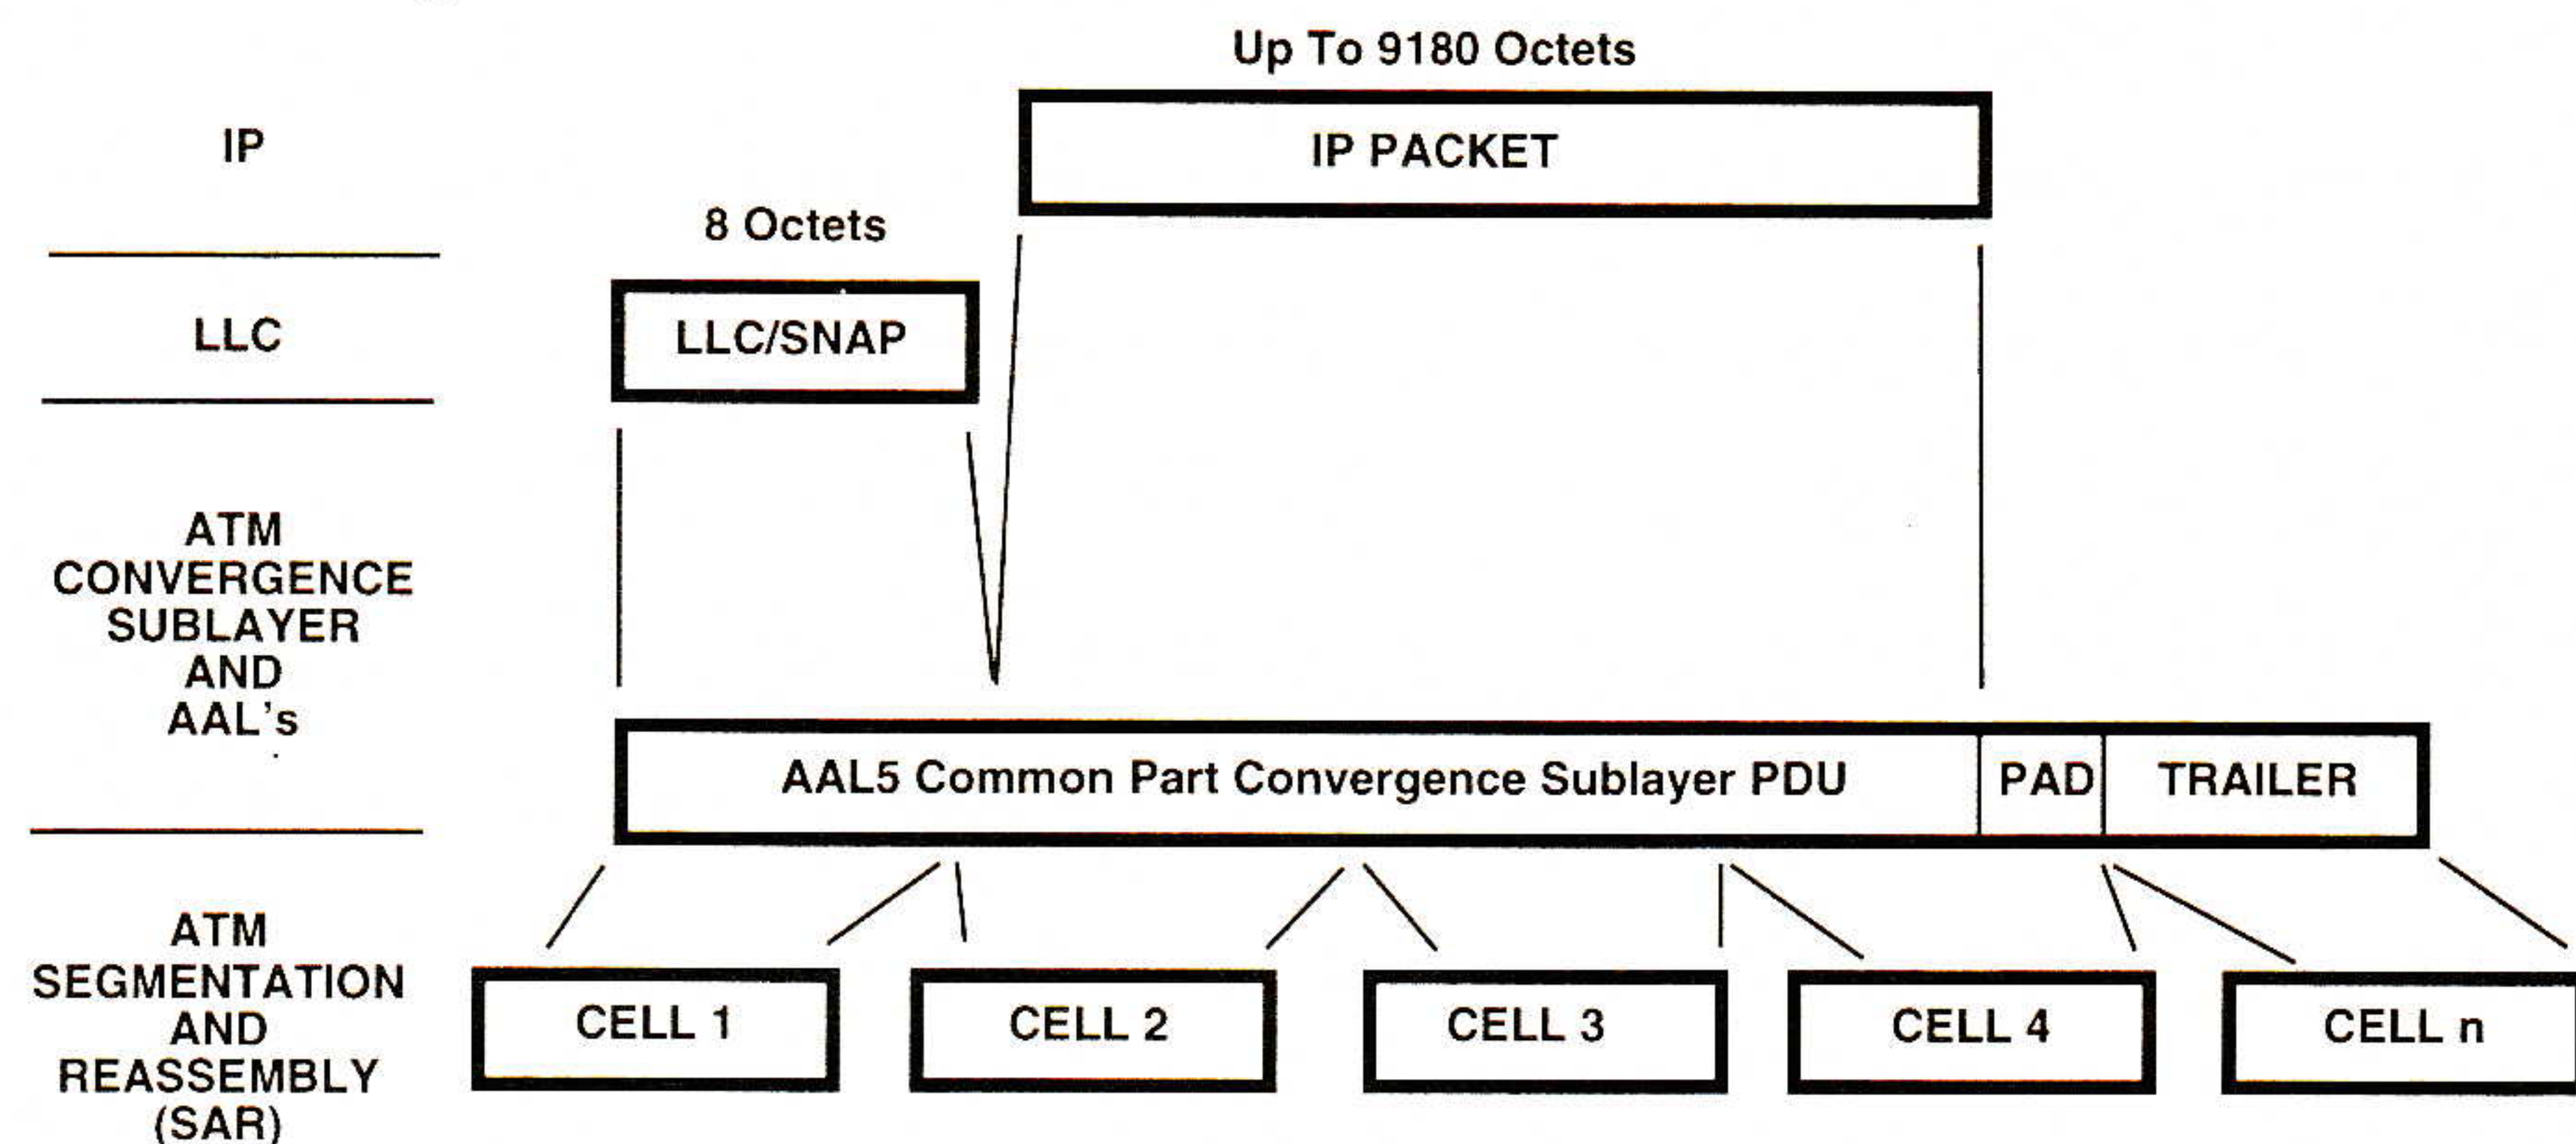


Figure 2: IP over ATM LLC/SNAP Encapsulation

Classical IP and ARP over ATM

Given a new data link layer, the fastest way to get IP running over it is to 1) specify the encapsulation (done in previous section) and 2) implement a model that looks, feels, and smells like IP subnets running over an Ethernet network. IP over ATM is kind of like IP over Ethernet/802.3 (looking from the top down) with some differences. First the similarities between the two approaches:

- The IP networking stacks sitting above the network interface and device driver look pretty much the same.
- An IP address is assigned to the ATM network interface just like assigning an IP address to an Ethernet LAN interface. That IP address, combined with the interface's subnetwork mask, defines the logical IP subnet (LIS) that the interface is participating in.
- When transmitting an IP packet, the subnet mask is applied to the destination address and the result checked. If the address is on the subnet, an address resolution mechanism is employed to determine the media address of the station and the packet is passed directly to the destination. If the destination is outside the local LIS, a routing function is employed to determine the next hop gateway for that packet, then the address resolution function is employed to deliver that packet to the next hop gateway on the local LIS.

So from an IP perspective, ATM and Ethernet/802.3 look pretty much the same with regards to the forwarding of packets. The differences between IP over ATM and IP over Ethernet/802.3 are:

- Ethernet/802.3 is a shared media access network: i.e., all stations share the same "cable," each station can listen, do collision detection, etc. ATM is a non-shared network. Each station has its own private attachment to the local area network. The switch fabric is the network.
- Ethernet/802.3 is inherently a multicast media that scales to the limits of the Ethernet specification (500 meters x 3, repeaters, bridges, etc.). ATM is not a multicast media. Limited multicast will be provided in the future via multicast service addressing and multicast servers. ATM local area networks can potential scale to very large sizes.

IP Over ATM & High-Speed Backbones (*continued*)

- In Ethernet/802.3 networks, frames are transmitted over the media with one IP packet typically per frame and each frame has a source and destination address. In ATM, cells are transmitted over the media with one IP packet being segmented into many cells with the restriction that each transmitting station has its own sending channel as seen by the receiver. For example, you don't want the cell streams from Stations A and B being mixed into a single receiving channel at Station C; ATM cells do not contain a source address. If this were the case, the cells for the packets from A and B would intermix and C would not be able to reassemble the individual packets.
- In Ethernet/802.3 networks, when a station wants to transmit to another station, it just waits for the shared media to become free and then transmits the frame. In ATM, a virtual channel must be set up first between the transmitting and receiving IP stations. ATM has *Permanent Virtual Channels* (PVCs) and dynamic or *Switched Virtual Channels* (SVCs). Permanent channels are set up by hand by the network administrators, with the caveat that all IP stations within an LIS must be connected to one another, requiring a full PVC mesh. Switched virtual channels are opened and closed as needed in an ATM network that supports a common signaling protocol [2]. After the VC is established, each station must agree as to which IP encapsulation method is being used: LLC/SNAP, null, or something else [3].
- In Ethernet/802.3 networks, the *Address Resolution Protocol* (ARP) is a broadcast protocol in which all stations participate; e.g., a station transmits an ARP query to all stations on the Ethernet segment and the responsible station then answers. In ATM PVC networks, the address resolution is performed via table lookup. The ARP table is constructed by hand as PVCs are configured into each system: i.e., when a PVC is created, the administrator says map this IP address to this PVC number. In ATM SVC networks, address resolution is performed via an ATMARP server mechanism [4]. This server resides within the LIS and each IP station must register with the server. When a station wants to transmit an IP packet, if its ATMARP table does not have the needed cached entry, a ARP query is sent to the ATMARP server, who then answers the query. In the future when ATM networks provide a general, universally deployed multicast mechanism, the ATMARP service may be replaced with a multicast service address and each IP station will participate in the ARP protocol similar to broadcast ARP for Ethernets.

Hopefully you can see that PVCs can be a real pain to configure if you have many hosts in your logical IP subnet.

Switched virtual circuits (SVCs) make it less of a pain to configure IP over ATM subnets, however in order to make the system work, an ATMARP server must be provided for the LIS and ATM Forum UNI 3.0 signaling [2] must be deployed on every IP ATM end point member in the LIS. Once these two system pieces are in place, subnets may be constructed over arbitrary ATM network topologies. For example, on one ATM switch two logical IP subnets can be constructed (A and B), each requiring its own ATMARP server. A and B each are different IP subnets (see Figure 3).

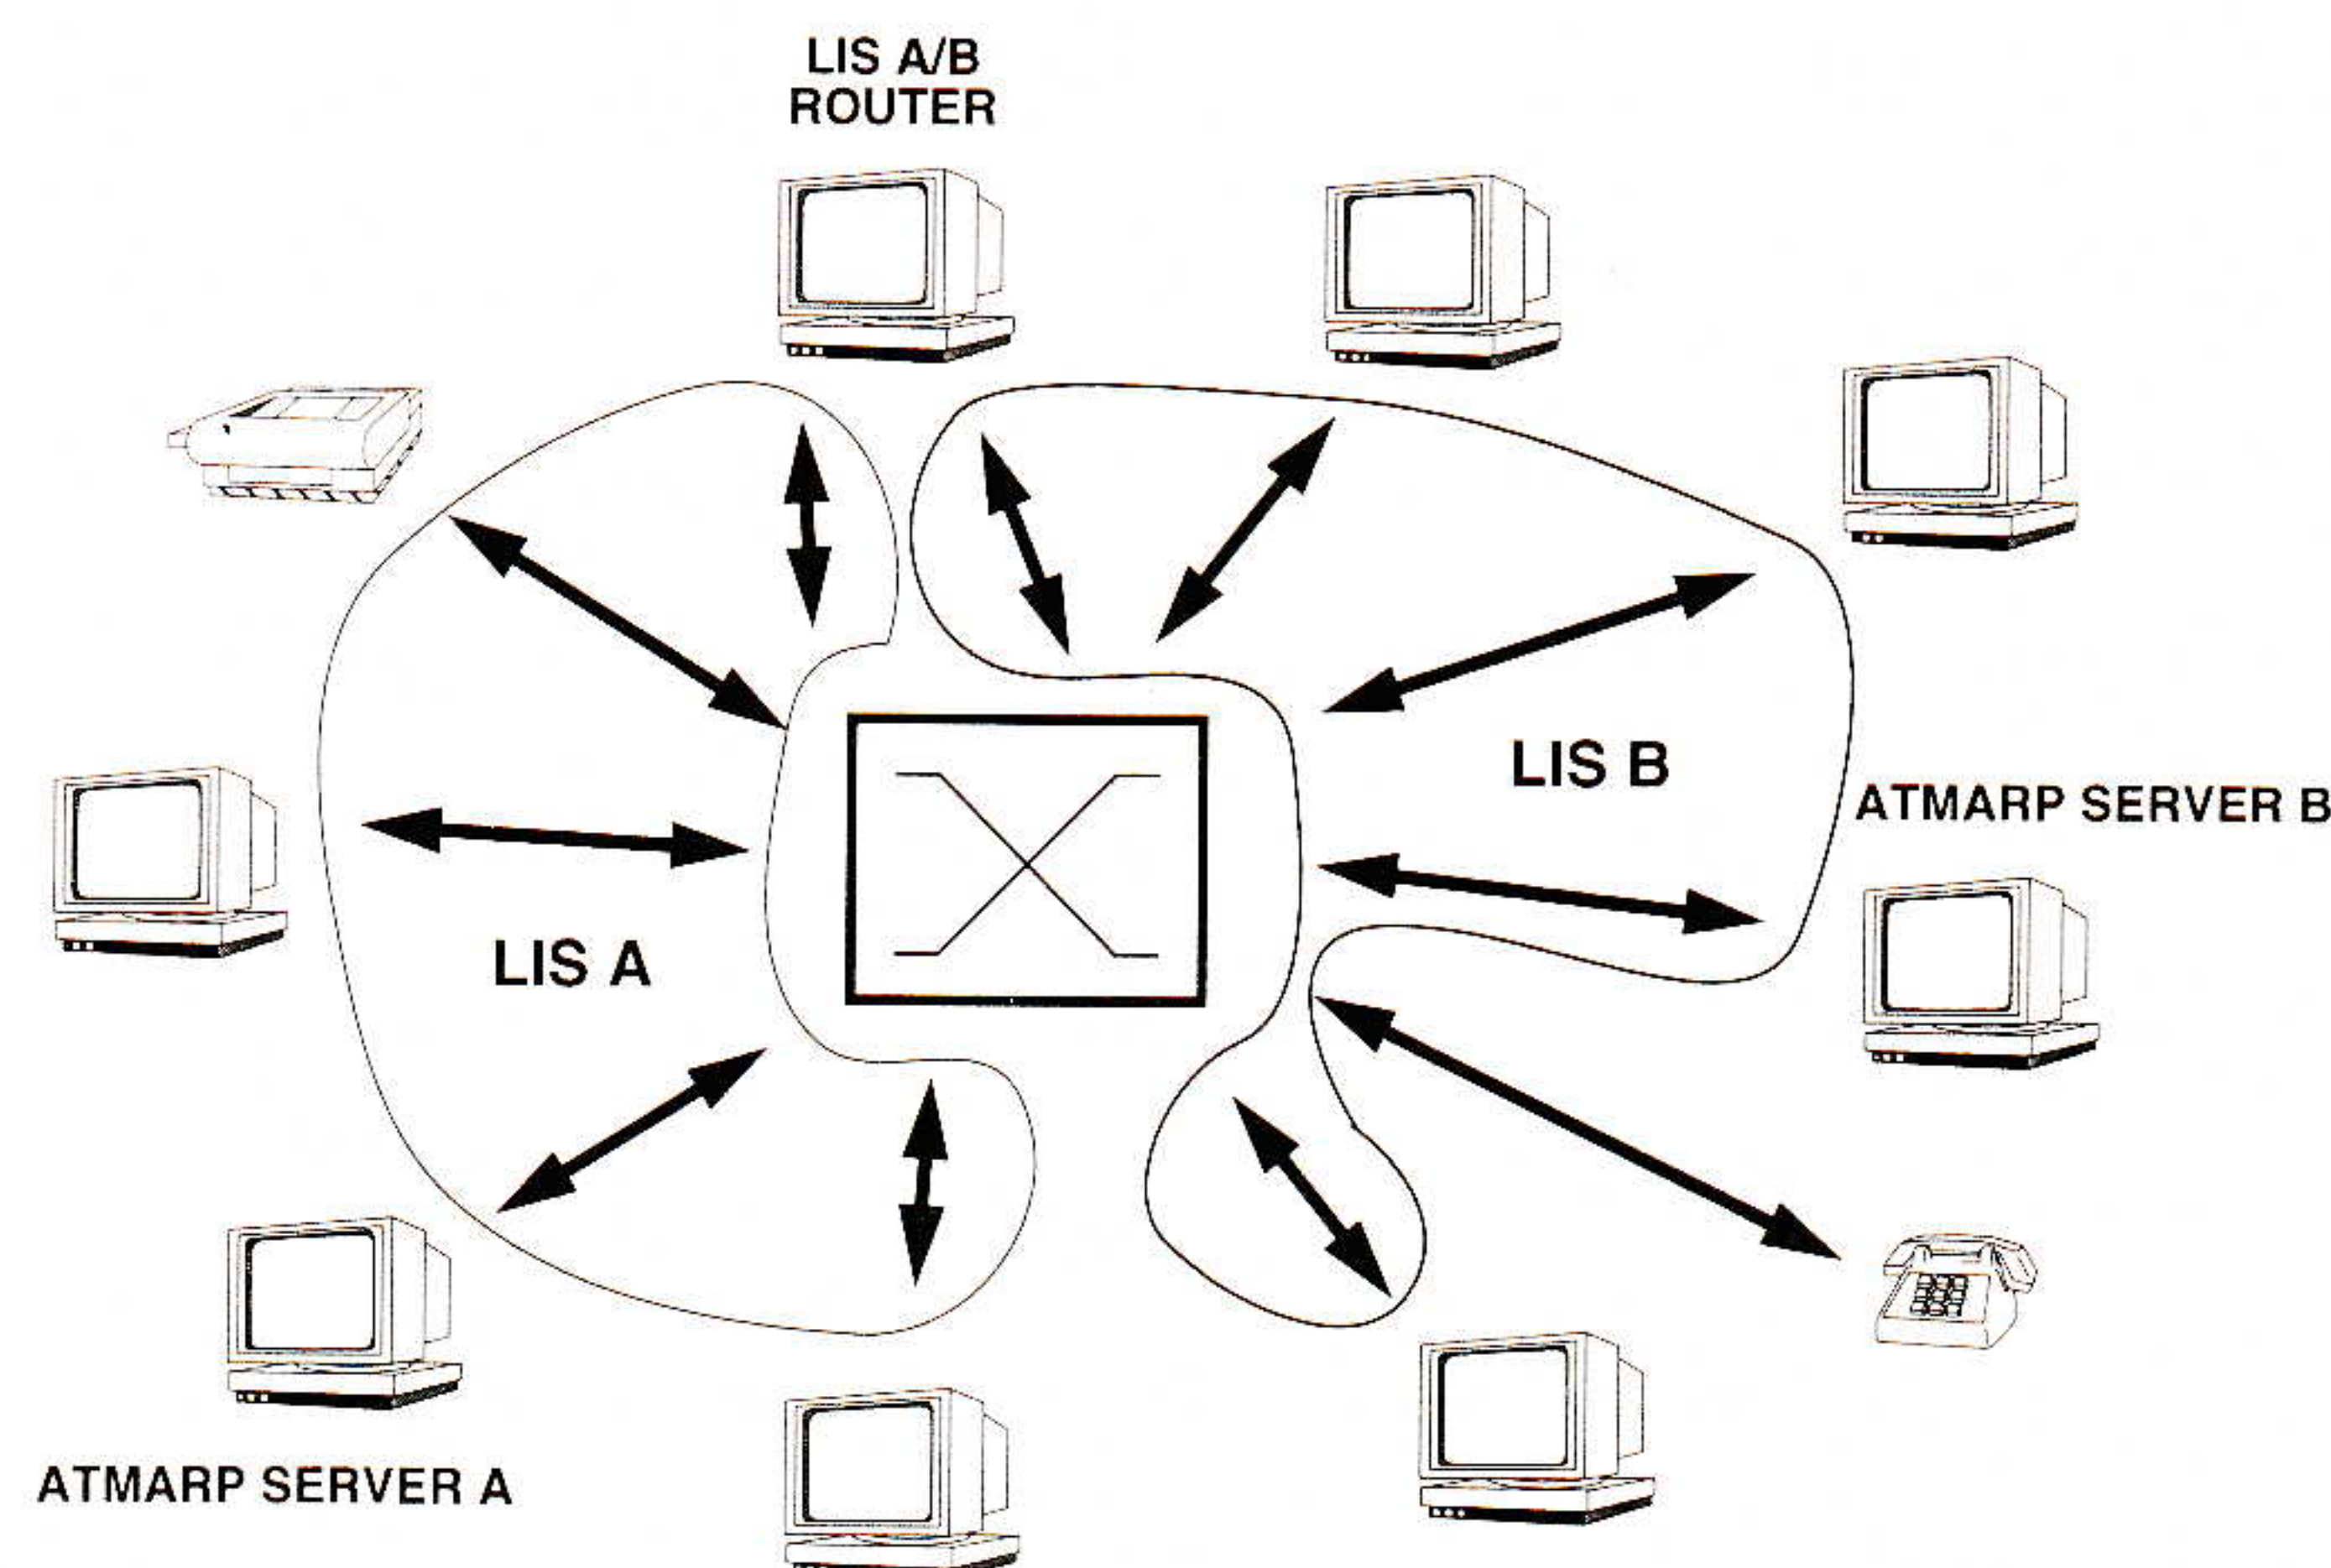


Figure 3. Logical IP subnets over ATM

When a station wishes to exchange IP packets with another station on the same LIS, a direct ATM virtual channel will be established between the two using the Q.93B signaling protocol [2]. When a station on A wishes to communicate with a station outside the subnet, say on subnet B, the sending station will set up a virtual channel to an IP router station on subnet A. The router station will then forward the IP packets as needed to reach subnet B. The Classical IP over ATM model as defined in RFC 1577 [4] requires an IP router function to forward IP packets between stations on different subnets.

Fourteen organizations are participating together to construct a nationally recognized gigabit ATM testbed in the San Francisco Bay Area. Pacific*Bell is the telecommunications provider for this network. Funding should be provided by Pacific*Bell's recently established *California Research and Education Network* (CALREN) program and by other government funding. BAGNET will be a two year project investigating a teleseminar application built over Pacific* Bell's production ATM network services offering.

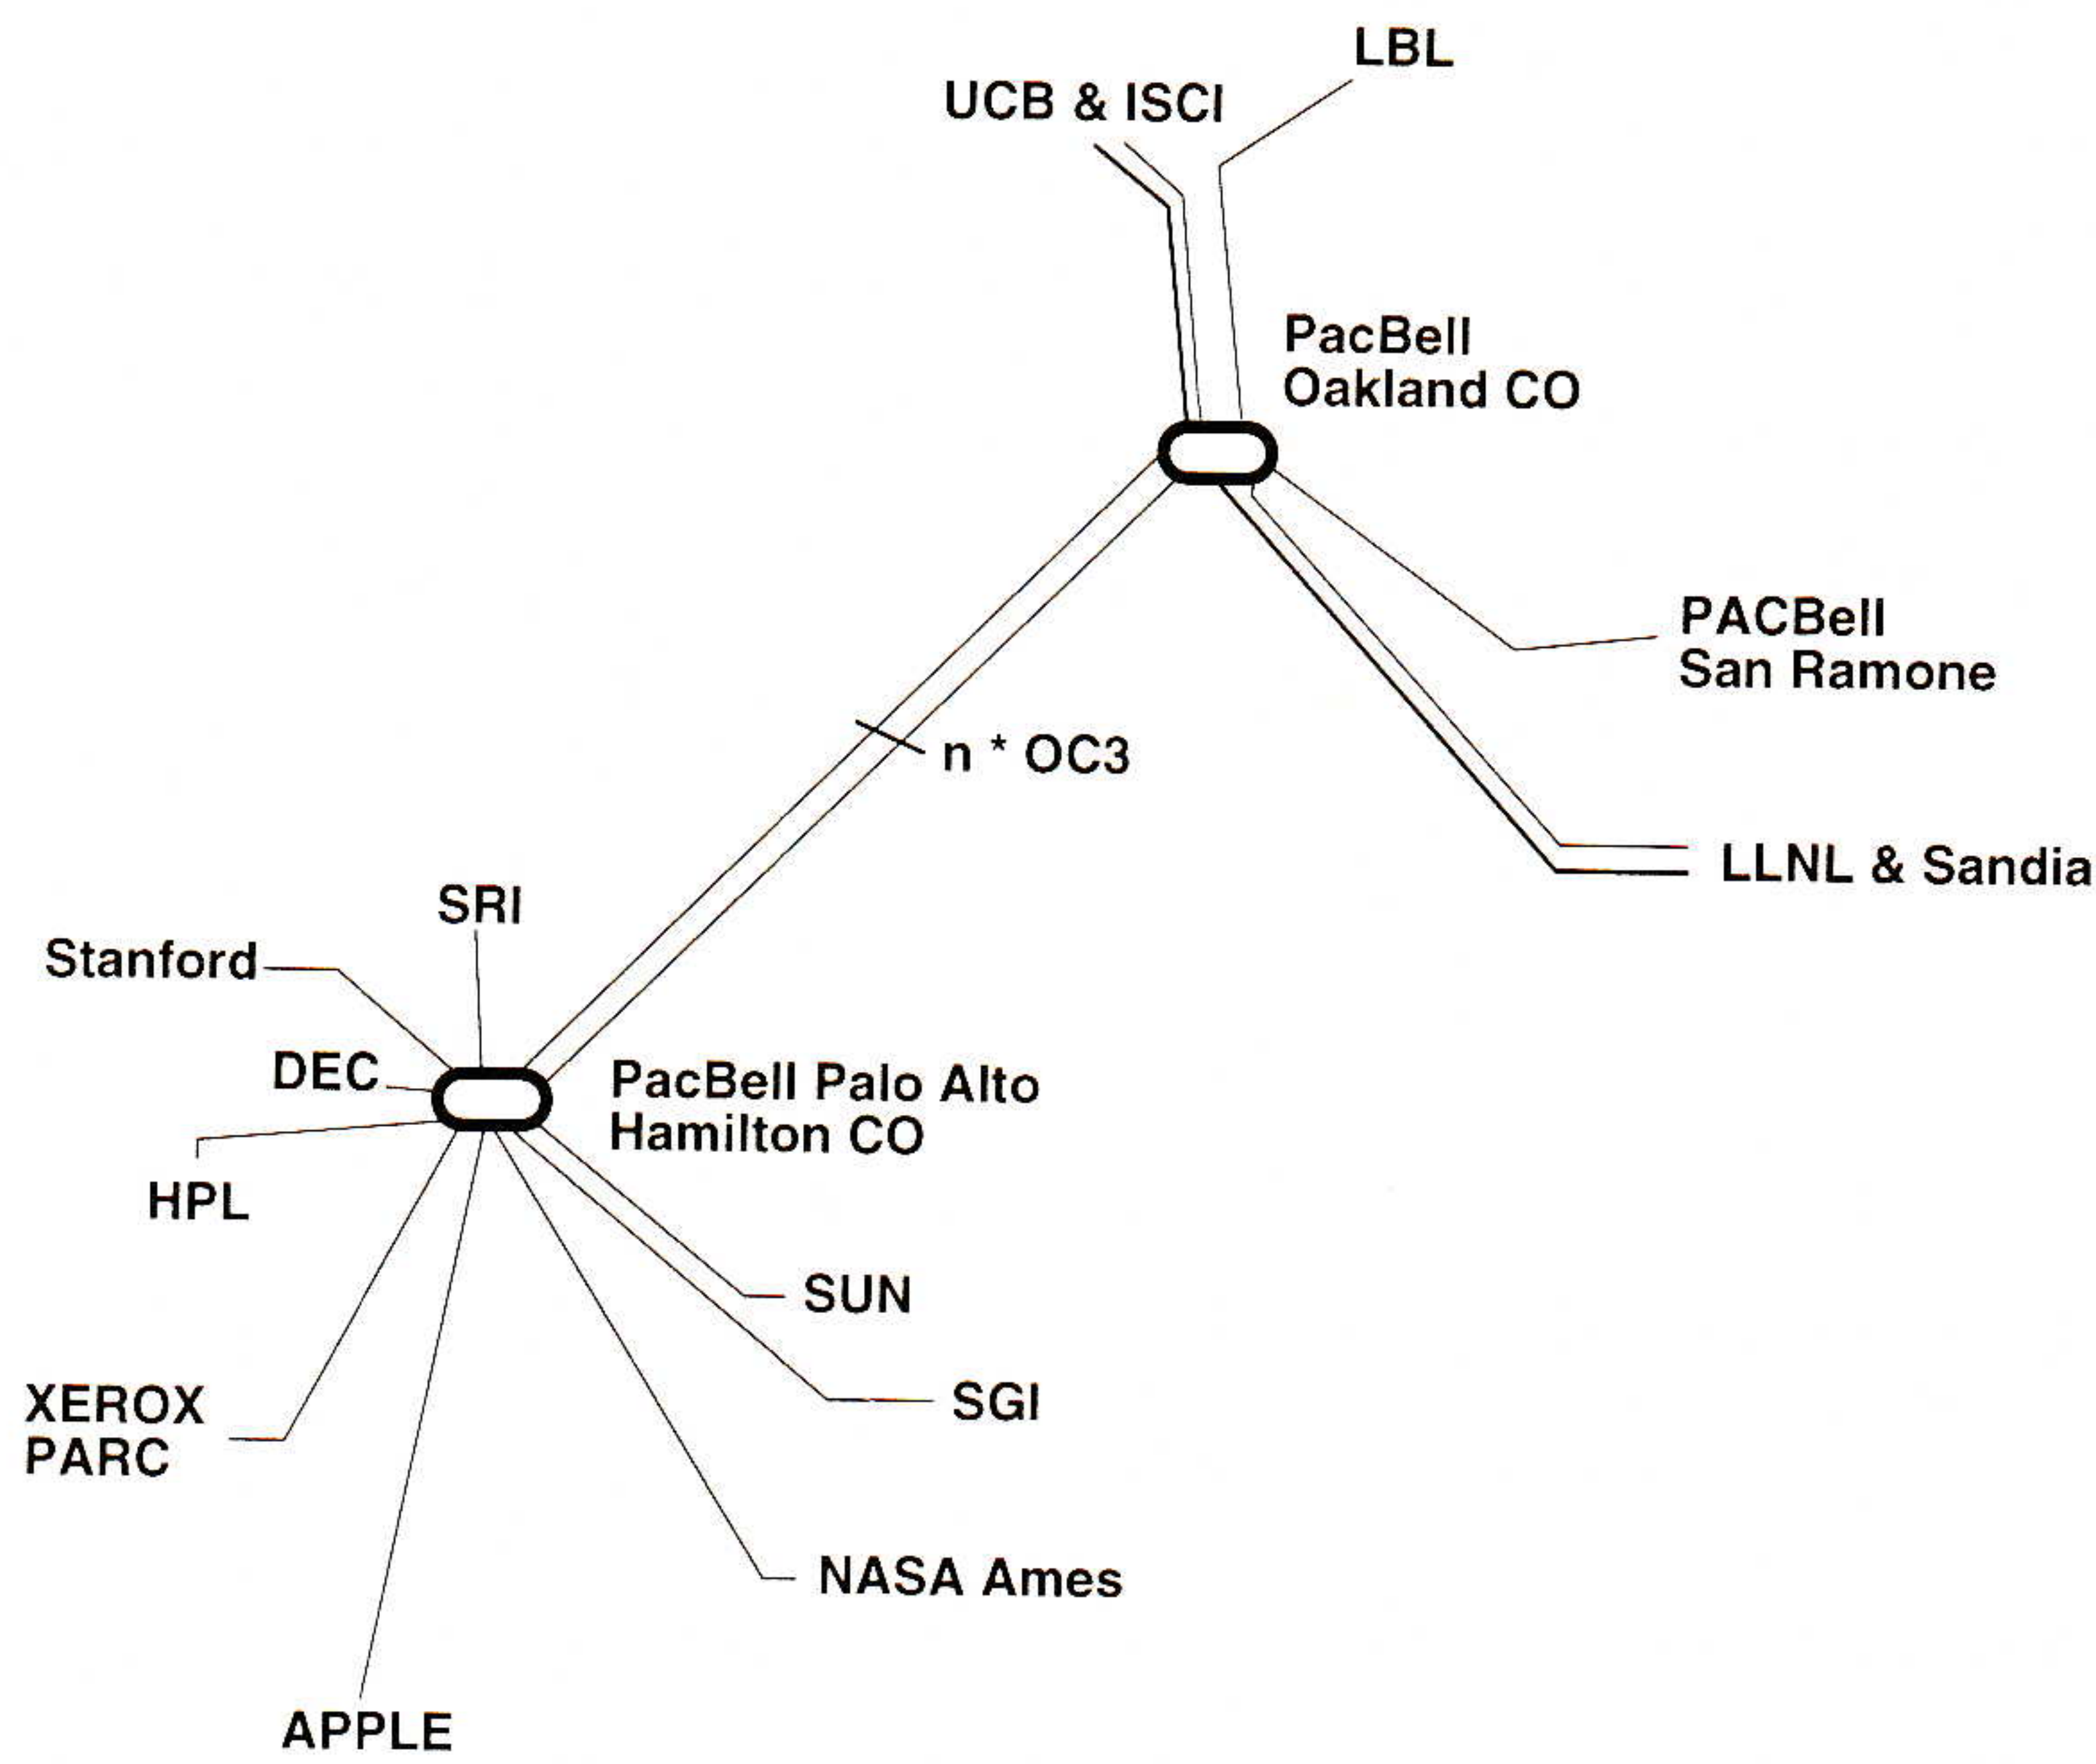


Figure 4: Bay Area Gigabit Testbed Topology

The Bay Area Gigabit Testbed (BAGNET)

The BAGNET members are: Apple Computer, DEC, Lawrence Berkeley Labs (LBL), Lawrence Livermore National Labs (LLNL), HP, ICSI, UC Berkeley, NASA Ames, Sandia, SGI, SRI International, Stanford University, Sun Microsystems, XEROX PARC, and Pacific*Bell.

IP Over ATM & High-Speed Backbones (*continued*)

Pacific*Bell formally turned on BAGNET as part of a marketing trial on December 30th, 1993. The first two sites connected were Xerox PARC and NASA Ames. The remainder of the sites were connected by the end of May, 1994. See Figure 4. Pacific*Bell will be providing production quality services and support for the network, in much the same way they provide their other tariffed services (voice, ISDN, etc.) There will be little room for experimentation with Pacific*Bell's ATM services.

Engineering BAGNET

Engineering usually implies the creation of a satisfactory solution, crafted with available technology, while obeying certain rules or constraints. We have a fine collection of constraints within BAGNET, some of these are technology imposed, some self-imposed. They are:

- Pacific*Bell has selected a vendor whose ATM switches:
 - Will not support Q.93B signaling (SVCs) until early 1995 [2].
 - Provides only 512 configurable VC entries per port.
 - Each host in the IP backbone subnet, must be fully connected PVC-wise to all other hosts on the backbone.
- If each site has 6 IP hosts and there are 15 sites, there will be 90 total hosts on the backbone. Each host must be connected to each other host outside the site. This requires roughly:
$$6 \text{ hosts} \times (6 \text{ hosts} \times \text{the 14 other sites}) = 504 \text{ full duplex PVCs}$$

...behind each site port, which is uncomfortably close to the 512 VC per site port limit allowing no headroom.
- The interconnect trunks between the Oakland and Palo Alto ATM switches must be few, due to cost reasons, yet support enough PVCs and bandwidth for the network. We figure on two trunks minimum connecting the two switches.

Given these constraints and 448 PVCs per port (512–64 for headroom), the number of PVCs will be the product of the number of hosts attached via the Palo Alto switch times the number of hosts attached via the Oakland switch.

- Multicast PVCs will be available on a limited basis, however most available host software will not initially be able to use it.

During the course of our *IP Down Under* planning meetings, we came to the following sets of design goals and constraints:

- Pacific*Bell will initially configure the complete BAGBONE allotment of PVCs, assuming a preset allotment of 2, 3, or 4 hosts at individual sites, selected per site. Tools will be created to aid in the tedious task of configuring PVCs throughout the network.
- BAGNET will not be used as an alternative path for production Internet use; i.e., BAGNET will not be used to avoid paying Internet access/use fees.
- The BAGBONE will not be used as a transit network for non-BAGNET use.
- As operation experience dictates, we will schedule high bandwidth needs. High bandwidth will be defined later.
- No routing protocols on the BAGBONE (initially), i.e., we'll use static configurations.

- Each site will have the option of making SNMP available to the network. Some sites are using the ATM switch fabrics for other users and will not be making SNMP available.
- BAGNET will follow RFC 1577 and RFC 1483; i.e., LLC/SNAP encapsulation will be required on all BAGBONE virtual channels.
- BAGBONE IP address assignments will be coordinated the old fashion way, i.e., with manual host tables, followed shortly by DNS support.
- If a site needs more than 3 or 4 hosts, then they will need to implement a local IP routing solution.
- Every PVC will be configured for the full link bandwidth (i.e., 155 Mbps). Our traffic management will be best effort with no peak limiting scheduling.
- We would like Pacific*Bell to urge its switch vendor to implement Q.93B signaling as soon as possible. PVCs are far too painful.

Our application experiments will consist of:

- Initially, getting the network up and running.
- Running the Internet Multicast Backbone (MBONE) tools over the BAGBONE, including mapping Class D IP addresses to point-to-multipoint VCs.
- Developing our motivating teleseminar application.

At the time of this writing, the current BAGNET installed base consists of XEROX PARC, NASA Ames, and HP Labs. We expect the other site to come on line by the end of May 1994. We have achieved limited operational testing at this time and are expecting significant results by the end of Summer 1994.

World Wide Web information

Information on the IP over ATM Working Group of the IETF can be found on the *World Wide Web* home page:

<http://matmos.hpl.hp.com/>

Information on BAGNET can also be found on the Web at:

<http://george.lbl.gov/BAGNet.html>

References

- [1] Laubach, Mark, "ATM for your internet—But When?" *Con-neXions*, Volume 7, No. 9, September, 1993.
- [2] ATM Forum, "User Network Interface (UNI) Specification Version 3.0," ISBN 0-13-225863-3, Prentice Hall, December 1993.
- [3] Heinänen, Juha, "Multiprotocol Encapsulation over ATM," RFC 1483, March, 1993.
- [4] Laubach, Mark, "Classical IP and ARP over ATM," RFC 1577, January, 1994.
- [5] Atkinson, Ran, "Towards Real ATM Interoperability," *Con-neXions*, Volume 7, No. 8, August 1993.

MARK LAUBACH holds a BEE. and M.Sc. from the University of Delaware. He has recently moved to Com21 Inc., a new venture in Mountain View, CA. Mark was with the Hewlett-Packard Company for 14 years and involved in a number of internet, ATM, and multimedia activities during his tenure there. Mark is a member of the IETF and is the chair of the IP-over-ATM Working Group. He is the author of RFC 1577 "Classical IP and ARP Over ATM." He is a past member of the CSNET Executive Committee and a past member of the board of trustees for the Corporation for Research and Educational Networking (CREN, CSNET+BITNET). This article was compiled when Mark was at the Hewlett-Packard Laboratories. His e-mail address is: laubach@netcom.com.

Book Reviews

The Internet Guide for New Users, by Daniel P. Dern, McGraw-Hill, Inc., 1994, ISBN 0-07-016511-4.

Audience

This book promises to demystify and explain the Internet for “anyone who has any type of computer and modem, a telephone, and budget of a dollar a day.” Subject to a few reservations, it succeeds. Mostly self-contained chapters allow reading selected material and skipping that which is of less interest.

Part 1 covers “ramping up, getting started.” Its 116 pages describe the Internet past and present, identify ways to gain access, provide network architectural concepts such as addressing and naming, and teach UNIX survival skills. UNIX, though widely used, is not universal. The book slights other platforms with its UNIX-centric perspective and examples. Though users are encouraged to investigate availability of Internet access on their work systems, the variety of environments they may encounter is not adequately addressed. Users with Net access are advised (or at least invited) to skip to Part 2.

Part 2 describes “The four basic Internet food groups”: e-mail, USENET, remote login, and FTP. This accurately indicates the critical nature of this section, which describes the key Internet application—the reasons people use the Net, and read this book.

Part 3, “Navigating the Net,” introduces Gopher,archie, WAIS, and similar applications, described as making “it much easier to use the Internet and make being on the Internet vastly more valuable.”

Part 4 deals with Internet citizenship, describing Net economics and access to human and online resources. Concepts such as network security, privacy, hygiene and manners are described, to help new users avoid pitfalls and rapidly become full-fledged network citizens.

Part 5 includes topics such as commercial services, archives, communities, and miscellany. E-mail lists, covered here, would be more logically found in Part 2, with basic email and USENET, since they are logical extensions of those facilities, rather than a unique Net community. As increasing numbers of Internet users arrive without benefit of organizational affiliation and support, more reliance is put on ad hoc resources such as user groups. These deserve more than the two sentences “User groups often offer help and support. Try your local computer society, etc.”

The automotive metaphor is frequently applied to the Internet and the (presumably) even more general Information Superhighway. In that context, this book is a good “tour guide,” illustrating highways, scenic routes, and tourist attractions, and even warning about the odd speed trap. It does not attempt to be a road map, leaving as an exercise for the driver acquiring complete site- or application-specific information and documentation.

Characters

The book quotes and describes many Internet players, from well-known Net names to Internet service providers to average users. These quotations illustrate how widespread and accessible the Internet is. Many quotee’s e-mail addresses are given, presumably indicating their willingness to receive mail from readers. These addresses can be a “starter set” of correspondents for new users wondering with whom that can use the Internet to communicate!

Chapters begin with appropriate and entertaining quotes from show business (*How to Succeed in Business...*, *The Grateful Dead*, *Gilbert and Sullivan*, *The Firesign Theatre*), literature (*Sherlock Holmes*, *Isaac Asimov*, *Alice in Wonderland*), and even the Net itself (`rec.humor.funny`):

*For everything, trn, trn, trn
There is a Newsgroup, trn, trn, trn*

And a thread for every subject under heaven Scenery There are many preferred learning styles. Some people devour reference manuals and some never touch them. Some people prefer descriptions of technology and others prefer illustrations of its use. This book is stronger for readers preferring narrative to examples, sometimes seeming (to me) to belabor background information before illustrating the application at hand. For example, 30 pages into the USENET chapter, one reaches the section "Start Your Newsreader Program," which begins "First, of course, you start up your newsreader program...." I would have preferred to start it 20 pages earlier, interleaving basic concepts with usage examples as the chapter proceeded.

Numerous boxed "sidebars" provide excellent information nuggets in the form of self-contained factoids or topics. These are directly useful pieces of advice or illustrations of concepts described in the main text. These would be more useful if included in the Table of Contents, or even in a separate easily accessed list.

A long, complete, and interesting Table of Contents invites both sequential reading and browsing. A comprehensive 26 page Index allows quickly locating topics of interest, or relating topical threads in different parts of the book.

The book does well at promising what will be told, telling, and then summarizing. Useful additions would be chapter-specific Network exercises to apply and reinforce information conveyed, and an overall part-specific quiz to illustrate key facts. The quiz might take the form of "The Internet Driver's Test" as printed in *Internet World* magazine, edited by the book's author.

Script

The book needs tighter organization, with elimination of repeated material. Even with readers advised to navigate the book according to their individual needs, there's no need to repeat (for example) descriptions of e-mail functions or the need for unique Internet resource names, each within two pages. Similarly, descriptions of USENET etiquette and general netiquette are each presented as multiple nearly adjacent lists. These should be consolidated for easier reader digestion.

Editing should have remedied frequent phrasing such as "There's the private networks...", and repaired the description of revealing one's e-mail inexperience as being an "e-mail newfie." Technical review should have ensured correct examples: the wrong value for shell variable \$user is shown, "/usr/alice" rather than "alice"; 7 bits equals 128 characters rather than 255; groups of 8 bits are more commonly called "bytes" or "characters" than "words"; and the UNIX pipe character is the vertical bar ("|") rather than the slash ("/") sometimes used.

Book Reviews (continued)**Aftermath**

Having once been rebuffed when suggesting that a colleague pursue Internet access with “But who would I talk to?” I am sensitive to the need to explain the “why” of the Net in addition to the “how.” Many individuals who have the book’s listed prerequisites (computer/modem/telephone/budget) still aren’t quite sure what the Net offers them. The book’s back cover and Preface provide lists of what this book will let readers do. But they both assume that potential readers know what the Net offers and understand its benefits. The Net’s “why” surely emerges from the body of the book. But addressing the “why” on the cover and in the Preface, by including benefits in addition to features (technology and applications) would greatly enlarge the potential readership.

It’s easy to forget how mysterious the Net can look to non-users, computer literate though they may be. Noting that Net access can make permanent positive changes in one’s business and even personal/social life, and giving examples of how e-mail and other Net applications can be a powerful hybrid technology combining phone, fax, radio, and newspaper, might greatly increase interest in climbing the learning curve of Net education, using this book as a tool.

—Gabriel Goldberg
gabe@access.digex.net

TCP/IP Illustrated, Volume 1: The Protocols, by W. Richard Stevens, Addison-Wesley, 1994, ISBN 0-201-63346-9.

The author of *TCP/IP Illustrated* has succeeded in creating another indispensable tome of networking knowledge. This is the most comprehensible and complete book I have read on TCP/IP. It takes a different slant than other books, by presenting not only details of TCP, IP, ARP, ICMP, routing, etc., but actually shows these protocols (and common Internet tools) in action.

Audience and organization

The book does not assume prior networking knowledge. The writing style is very readable, and it gives good clear introductions to all subjects without coddling the reader. I think it is well aimed for someone taking their first look at TCP/IP, as well as people who have done some client-server programming, and want to know how everything really works. Each chapter ends with a summary section and questions to make sure key points are understood. Answers to the questions are located at the back of the book.

The book begins with a gentle introduction to TCP/IP and protocol layering. The rest of the book follows a bottom up approach, beginning with the basic protocols going up through details on FTP, SMTP, and NFS. The Appendices describe diagnostic tools that are used in the text, and sources that they can be FTP’d from.

It explains the material by first giving the standard datagram layouts and a description of the various fields and how they are to be used. It then differs from other texts by using publically available network diagnostic tools (largely *tcpdump*) to show the interactions that are actually taking place over the network.

This is where the “illustrated” part comes in. In illustrating these interactions, the author makes use of a sample network of heterogeneous hosts running over an Ethernet, over SLIP, and across a router. Differences in protocol implementations on the hosts are highlighted in the text. The inside front cover contains a map of the sample network for easy referencing, and the inside back cover has an equally useful acronym dictionary—which greatly helps with choking down the alphabet soup endemic to networking literature.

Practical perspective

One very refreshing aspect of this book is its *practical* perspective. The protocols are presented according to the standard, but the author is very clear in pointing out where certain implementations differ from the standard.

Demonstrating the protocols over a realistic network, and using publically available tools to monitor the protocols is another example of the books practicality.

Highly recommended

TCP/IP Illustrated emphasizes conceptual understanding of how the protocols actually work. After reading the book I gained both a better understanding of how TCP/IP works, and how I might use network diagnostic tools to locate a problem. I highly recommend this book to anyone interested in learning about TCP/IP.

—Eli Charne, UCI


Write to *ConneXions*!

We'd love to hear your comments, suggestions and questions about anything you read in *ConneXions*. Our editorial address is given below. Use this address for letters to the Editor, questions about back issues etc.:

ConneXions—The Interoperability Report
303 Vintage Park Drive
Suite 201
Foster City, CA 94404-1138
USA
Phone: +1 415-578-6900 or 1-800-INTEROP (Toll-free in the USA)
Fax: +1 415-525-0194
E-mail: connexions@interop.com

For questions about your subscription please call our customer service hotline: 1-800-575-5717, (+1 502-493-3217 outside the US). This is the number for our subscription agency, the Cobb Group. Their fax number is +1 502-491-8050. The mailing address for subscription payments is P.O. Box 35840, Louisville, KY 40232-9496.

NetWorld+Interop World Tour



Atlanta, GA:	September 12–16, 1994
Paris, France:	October 24–28, 1994
Las Vegas, NV:	March 27–31, 1995
Berlin, Germany:	June 26–30, 1995
Tokyo, Japan:	July 17–21, 1995

See you there!

Announcement and Call for Submissions

The *USENIX Winter 1995 Technical Conference* will be held January 16–20, 1995 in New Orleans, and will be the only broad-theme USENIX conference in 1995. The emphasis for the USENIX Winter 1995 Conference is on state-of-the-art practice and research in personal, distributed, and enterprise computing.

Topics

We seek original and innovative papers about the architecture and performance of modern computing systems. We are especially interested to hear reports on practical experiences with such systems. Of particular interest are such topics as:

- Privacy and cryptography
- Personal Digital Assistant applications
- Enterprise-scale computing
- Kernelized operating systems
- User interface toolkits
- Standards-based computing environments
- File systems and mass storage
- Nomadic and wireless computing
- Shared address spaces

The USENIX conference, like most conferences and journals, requires that papers not be submitted simultaneously to more than one conference or publication and that submitted papers not be previously or subsequently published elsewhere. Papers accompanied by so-called “non-disclosure agreement” forms are not acceptable and will be returned to the author(s) unread. All submissions are held in the highest confidentiality prior to publication in the Proceedings, both as a matter of policy and in accord with the U.S. Copyright Act of 1976 (Title 17, U.S. Code, Section 102).

Submission guidelines

It is important that you contact the USENIX Association office to receive detailed guidelines for submitting a paper to the refereed track of the technical sessions; please telephone to +1-510-528-8649 or E-mail to winter95authors@usenix.org. In addition, specific questions about submissions to the USENIX Winter 1995 Conference may be made to the program chair via Internet e-mail at honey@citi.umich.edu.

The program committee will review full papers or extended abstracts. An extended abstract should be 5 manuscript pages (single-sided) or fewer in length. It should represent the paper in “short form.” Please include the abstract as it will appear in the final paper. If the full paper has been completed, it may be submitted instead of an extended abstract. Full papers should be limited to 12 single-spaced pages.

Include references to establish that you are familiar with related work, and, where possible, provide detailed performance data to establish that you have a working implementation and measurement tools.

Every submission should include one additional page or separate e-mail message containing:

- The name of one of the authors, who will act as the contact for the program committee
- Contact's surface mail address, daytime and evening telephone numbers, e-mail address, and fax number
- An indication of which, if any, of the authors are full-time students

Where to send submissions

Submit one copy of an extended abstract or full paper by July 18, 1994 via *at least two* of the following methods:

- E-mail to winter95papers@usenix.org
- FAX to +1 313-763-4434
- Mail to:

Winter 1995 USENIX
CITI
University of Michigan
519 W. William
Ann Arbor, MI 48103-4943
USA

Important dates

Manuscripts or Extended Abstracts Due: July 18, 1994
Notification to Authors: August 31, 1994
Camera-ready Papers Due: November 14, 1994

Cash prizes

Cash prizes will be awarded for the best paper at the conference and the best paper by a full-time student.

Program Committee

Charles J. Antonelli	CITI, University of Michigan
David Bachmann	IBM Austin
David Chaum	DigiCash b.v.
Cecelia D'Oliviera	Information Systems, MIT
Richard Draves	Microsoft Research
Lori Grob	Chorus Systemes
Peter Honeyman (Chair)	CITI, University of Michigan
John T. Kohl	Atria Software
Greg Minshall	Novell, Inc.
Douglas Orr	University of Utah
Noemi Paciorek	Horizon Research
Phil Winterbottom	AT&T Bell Laboratories

More information

Materials containing all details of the technical sessions and tutorial program, conference registration, hotel discounts, and airfare discount and reservation information will be available at the end of September 1994. If you wish to receive the registration materials, please contact:

USENIX Conference Office
22672 Lambert St., Suite 613
Lake Forest, CA 92630
USA
Phone: +1 714-588-8649
Fax: +1 714-588-9706
E-mail: conference@usenix.org

CONNEXIONS

303 Vintage Park Drive
Suite 201
Foster City, CA 94404-1138
Phone: 415-578-6900
FAX: 415-525-0194

FIRST CLASS MAIL
U.S. POSTAGE
PAID
SAN JOSE, CA
PERMIT NO. 1

ADDRESS CORRECTION
REQUESTED

CONNEXIONS

EDITOR and PUBLISHER Ole J. Jacobsen

EDITORIAL ADVISORY BOARD Dr. Vinton G. Cerf
Senior Vice President, MCI Telecommunications
President, The Internet Society

A. Lyman Chapin, Chief Network Architect,
BBN Communications

Dr. David D. Clark, Senior Research Scientist,
Massachusetts Institute of Technology

Dr. David L. Mills, Professor,
University of Delaware

Dr. Jonathan B. Postel, Communications Division Director,
University of Southern California, Information Sciences Institute



Printed on recycled paper

Subscribe to CONNEXIONS

U.S./Canada ☐ \$150. for 12 issues/year ☐ \$270. for 24 issues/two years ☐ \$360. for 36 issues/three years

International \$ 50. additional per year (Please apply to all of the above.)

Name _____ Title _____

Company _____

Address _____

City _____ State _____ Zip _____

Country _____ Telephone () _____

☐ Check enclosed (in U.S. dollars made payable to CONNEXIONS).

☐ Visa ☐ MasterCard ☐ American Express ☐ Diners Club Card # _____ Exp. Date _____

Signature _____

Please return this application with payment to:

CONNEXIONS

Back issues available upon request \$15./each
Volume discounts available upon request

303 Vintage Park Drive, Suite 201
Foster City, CA 94404-1138
415-578-6900 FAX: 415-525-0194
connexions@interop.com

CONNEXIONS